

NAVAL POSTGRADUATE SCHOOL
Monterey, California



THESIS

**TRUST AND ITS RAMIFICATIONS FOR THE DOD
PUBLIC KEY INFRASTRUCTURE (PKI)**

by

Carl M. Pedersen

March 2001

Thesis Co-Advisors:

James Bret Michael
Audun Jøsang

Approved for public release; distribution is unlimited.

20010316 068

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

Public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instruction, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing this burden, to Washington headquarters Services, Directorate for Information Operations and Reports, 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302, and to the Office of Management and Budget, Paperwork Reduction Project (0704-0188) Washington DC 20503.

1. AGENCY USE ONLY (Leave blank)	2. REPORT DATE March 2001	3. REPORT TYPE AND DATES COVERED Master's Thesis	
4. TITLE AND SUBTITLE Trust and its Ramifications for the DoD Public Key Infrastructure (PKI)		5. FUNDING NUMBERS	
6. AUTHOR(S) Pedersen, Carl M.			
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Naval Postgraduate School Monterey, CA 93943-5000		8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)		10. SPONSORING / MONITORING AGENCY REPORT NUMBER	
11. SUPPLEMENTARY NOTES The views expressed in this thesis are those of the author and do not reflect the official policy or position of the Department of Defense or the U.S. Government.			
12a. DISTRIBUTION / AVAILABILITY STATEMENT Approved for public release; distribution is unlimited.		12b. DISTRIBUTION CODE	
13. ABSTRACT (maximum 200 words) Researchers have used a wide variety of trust definitions, leading to a plethora of meanings of the concept. But what does the word 'trust' mean? While most scholars provide their own definition of trust, they are dissatisfied regarding their own lack of consensus about what trust is. Trust is a cognitive function and modeling trust is an attempt to emulate the way a human assesses trust. Models of trust have been developed in an attempt to automate the logic, variables, and thought processes that a human performs when making a trust-decision. This thesis evaluates the various forms of trust and trust models. The results from our research found no such model that incorporates both mandatory and discretionary trust. A new hybrid model will be introduced, the "D-M Model." The motivation for using our model in the context of trust stems primarily from the appropriate use of discretionary and mandatory trust policies in organizations to ensure precision, consistency, and added assurance in trust. The real value of the D-M model, is that it addresses the need to model both of these types of policies explicitly and concurrently. This thesis concludes with the assessment of two practical applications of the D-M trust model as it is applied to DoD's Joint Task Forces.			
14. SUBJECT TERMS Trust models, Trust management, PKI, Computer Security		15. NUMBER OF PAGES 108	
		16. PRICE CODE	
17. SECURITY CLASSIFICATION OF REPORT Unclassified	18. SECURITY CLASSIFICATION OF THIS PAGE Unclassified	19. SECURITY CLASSIFI- CATION OF ABSTRACT Unclassified	20. LIMITATION OF ABSTRACT UL

NSN 7540-01-280-5500

Standard Form 298 (Rev. 2-89)
Prescribed by ANSI Std. Z39-18

Approved for public release; distribution is unlimited

**TRUST AND ITS RAMIFICATIONS FOR THE DOD PUBLIC KEY
INFRASTRUCTURE (PKI)**

Carl M. Pedersen
Lieutenant, United States Navy
B.S., Oregon State University, 1995

Submitted in partial fulfillment of the
requirements for the degrees of

MASTER OF SCIENCE IN INFORMATION SYSTEMS AND OPERATIONS

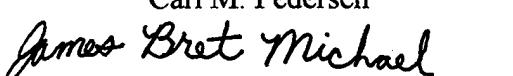
from the

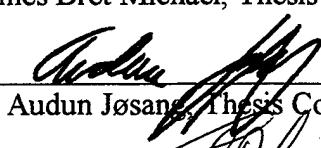
NAVAL POSTGRADUATE SCHOOL
March 2001

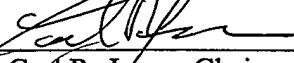
Author:


Carl M. Pedersen

Approved by:


James Bret Michael
James Bret Michael, Thesis Co-Advisor


Audun Jøsang, Thesis Co-Advisor


Carl R. Jones, Chairman
Information Systems and Operations
Curriculum Committee

ABSTRACT

Researchers have used a wide variety of trust definitions, leading to a plethora of meanings of the concept. But what does the word 'trust' mean? While most scholars provide their own definition of trust, they are dissatisfied regarding their own lack of consensus about what trust is. Trust is a cognitive function and modeling trust is an attempt to emulate the way a human assesses trust. Models of trust have been developed in an attempt to automate the logic, variables, and thought processes that a human performs when making a trust-decision. This thesis evaluates the various forms of trust and trust models. The results from our research found no such model that incorporates both mandatory and discretionary trust. A new hybrid model will be introduced, the "D-M Model." The motivation for using our model in the context of trust stems primarily from the appropriate use of discretionary and mandatory trust policies in organizations to ensure precision, consistency, and added assurance in trust. The real value of the D-M model, is that it addresses the need to model both of these types of policies explicitly and concurrently. This thesis concludes with the assessment of two practical applications of the D-M trust model as it is applied to DoD's Joint Task Forces.

TABLE OF CONTENTS

I. INTRODUCTION	1
A. AREA OF RESEARCH.....	1
B. SCOPE	1
C. RESEARCH GOAL.....	2
D. DEVELOPING TRUST.....	3
E. IS TRUST TRUSTWORTHY	8
F. TRUST IN THE INTERNET ENVIRONMENT	12
G. SECURITY CLEARANCE.....	15
H. TRUST IN DIGITAL SIGNATURE.....	17
II. BACKGROUND.....	19
A. TRUST	19
B. THE TOFFLERS AND THE WAVE OF TRUST	23
C. GENERAL SECURITY ISSUES	26
D. APPLICATIONS TO COMPUTER SECURITY	29
E. TRUST IN COMMERCIAL-OFF-THE-SHELF (COTS) PRODUCTS	30
III. MODELS	33
A. FORMAL MODELS	33
B. BASIC CONCEPTS TO BE MODELED	34
1. Strict Hierarchy	34
2. Distributed Trust Architecture	37
a. Mesh Configuration (Distributed Trust Architecture)	39
b. Hub-and-Spoke Configuration.....	39
3. Web Model.....	40
4. User-Centric Trust	42
C. RADIA PERLMAN	43
IV. D-M Model	53
A. ISSUES OF TRUST	53
B. TRUST SYSTEMS.....	54
C. D-M MODEL.....	55

D.	THE BASIC COMPONENTS OF THE D-M MODEL	56
1.	Actors	56
2.	Groups	56
3.	Levels Of Arbitration	57
4.	Policies	58
a.	Discretionary Policy.....	58
b.	Mandatory Policy	59
c.	Interaction Between Policies.....	59
E.	LIMITATIONS OF D-M MODEL.....	60
F.	SUMMARY	61
V.	CASE STUDY	63
A.	CASE STUDY - AEGIS PLATFORM.....	63
B.	CASE STUDY – BATTLE GROUP CONNECTIVITY.....	66
C.	TOOLS TO ASSIST NAVY SHIPBOARD DECISION MAKERS	69
1.	The Rapid Anti-Ship Cruise Missile Integrated Defense System....	69
2.	Ship Self-Defense System (SSDS).....	70
a.	The Principal Threat.....	70
b.	System Description	71
3.	Cooperative Engagement Capability.....	71
D.	SUMMARY OF KEY POINTS	73
VI.	CONCLUSION.....	75
A.	SUMMARY	75
B.	FUTURE WORK	77
1.	Public Key Infrastructure	77
2.	Influence Net Modeling	78
3.	Decision Support	79
	APPENDIX. GLOSSARY	81
	LIST OF REFERENCES	85
	INITIAL DISTRIBUTION LIST	89

LIST OF FIGURES

FIGURE 1, STRICT HIERARCHY (FROM ADAMS, 1999)	36
FIGURE 2, DISTRIBUTED TRUST ARCHITECTURE, (FROM ADAMS, 1999)	38
FIGURE 3, WEB MODEL, (FROM ADAMS, 1999).....	41
FIGURE 4, USER-CENTRIC MODEL, (FROM ADAMS, 1999)	43
FIGURE 5, A SINGLE – CA MODEL (FROM PEARLMAN, 1999)	44
FIGURE 6, A SINGLE – CA PLUS RA (FROM PEARLMAN, 1999)	45
FIGURE 7, AN OLIGARCHY OF CAs (FROM PEARLMAN, 1999).....	46
FIGURE 8, CONFIGURED PLUS DELEGATED CAs (FROM PEARLMAN, 1999)	47
FIGURE 9, ANARCHY (FROM PEARLMAN, 1999)	48
FIGURE 10, UP-CROSS-DOWN (FROM PEARLMAN, 1999).....	50
FIGURE 11, D-M MODEL	55

ACRONYMS

ANSI	American National Standards Institute
API	Application Program Interface
ASD(C3I)	Assistant Secretary of Defense, Command, Control, Communications and Intelligence
CA	Certification Authority
CGI	Common Gateway Interface
CMS	Communication Security Material System
CRL	Certificate Revocation List
DES	Data Encryption System
DISA	Defense Information Systems Agency
DNS	Domain Name Server
DoD	United States Department of Defense
DSL	Digital Subscriber Line
FIPS	Federal Information Processing Standard
GAO	Government Accounting Office
GUI	Graphical User Interface
IA	Information Assurance
IETF	Internet Engineering Task Force
IMAP	Internet Message Access Protocol
IP	Internet Protocol
ITSEC	Information Technology Security Evaluation Criteria
KA	Key Authenticity
LAN	Local Area Network
LDAP	Lightweight Directory Access Protocol
MAC	Message Authentication Code
MISSI	Multi-Level Information Systems Security Initiative
NATO	North Atlantic Treaty Organization
NIC	Network Interface Card
NSA	National Security Agency
OCSP	On-Line Certificate Status Protocol
OLAP	On-Line Analytical Processing
ORA	Organizational Registration Authority
PAA	Policy Approval Authority
PCA	Policy Creation Authority
PGP	Pretty Good Protection
PICS	Platform for Internet Content Selection
PIN	Personal Identification Number
PKCS	Public Key Certificate Standard
PKI	Public Key Infrastructure
PKIX	Public Key Infrastructure (X.509)
PMA	Policy Management Authority
POP	Post Office Protocol

RA	Registration Authority
RRQ	Recommendation Request Message
RSA	Rivest, Shamir, and Adleman
TA	Trusted Authority
TAO	Tactical Action Officer
TCB	Trusted Computing Base
TCSEC	Trusted Computer Security Evaluation Criteria
WWW	World Wide Web

ACKNOWLEDGEMENT

The author would like to thank Bret Michael for his invaluable assistance, guidance, encouragement, and editorial skill. His dedication to this thesis was above and beyond the call of duty. The author would also like to thank Audun Jøssang for his comments and encouragement. Without the help of these gentlemen, the thesis would not have been possible.

I. INTRODUCTION

A. AREA OF RESEARCH

This thesis consists of an investigation of the role of discretionary-mandatory trust models in today's computing paradigms for the protection of sensitive information. A hybrid model consisting of discretionary and mandatory policy regarding trust could create a new standard for how trust is represented and manipulated in a computational form. Currently, there are no models which accommodate both discretionary and mandatory policy about trust. The product of the research reported in this thesis is a model of discretionary-mandatory policy. There are some cases in which one would want discretionary control. For example: If the Naval Postgraduate School utilizes a Public Key Infrastructure (PKI) and student *A* has a relationship with student *B*, we would not want the Department of Defense to dictate the amount of trust between student *A* and student *B*. We would want to allow for some discretion. If however, the Naval Postgraduate School were to have a relationship with an outside university such as MIT, one would want to have mandatory policy in place, with the policy having been specified at the upper echelons of the chain of command.

B. SCOPE

This thesis builds upon the research of Gaines' (2000), in which he reviewed several trust models as they apply to the DoD Public Key Infrastructure. We have assessed additional models of trust and investigated the role of trust as it applies to the U.S. Navy. We define the concept of trust and both compare and contrast various trust models by evaluating their characteristics, environmental references, metrics, variables

used, and outputs. We also apply these concepts to a new hybrid model and demonstrate the models use in the context of network centric warfare and information operations.

One of the challenges in conducting this research is that the DoD's information infrastructure is constantly changing: it is a moving target. Thus, representation frameworks used to model policy regarding trust must be extensible to accommodate change.

C. RESEARCH GOAL

Trust is a cognitive function. Trust modeling is an attempt to emulate the way a human assesses trust. There are a number of trust models that represent attempts to define and assign metrics to trust. These models address the notion of trust in many different ways and their definitions and metrics vary significantly. A bewildering array of meanings and connotations of trust are available: there is no consensus on what trust means. In fact, if one examines the many definitions, one might come to the conclusion that existing trust models are an amalgamation of different beliefs and ideas.

Depending upon an entity's security policy and how that entity chooses to implement trust models, multiple levels of trust may have to be addressed. Additionally, interpretations of trust can differ among computing bases, domains, and applications. This makes it challenging to convey trust between entities on the Internet. In this thesis we assess existing modeling frameworks used to model trust. In the remainder of the thesis, we compare and contrast trust models by identifying their characteristics, environmental references, metrics, variables used, and outputs. The thesis

concludes with a case study of the requirements for modeling and reasoning about trust in a network centric warfare environment.

D. DEVELOPING TRUST

The effective use of information technology and success in any organization requires trust, not only of the information communicated, but also among faceless communicators. Our belief in the validity of the complex and subtle messages we receive by telephone or e-mail are conditioned on how well we know and trust the senders. In a sense, psychological bandwidth varies directly with the degree of trust between people. Trust cannot be decreed. The willingness to trust is a combination of values and evaluation, attitudes and interests. National culture influences how and whom we trust. But within and across cultures, trust depends on whom we consider trustworthy and how well we create trust in others. (Maccoby, 1997)

The ideas in Joint Vision (JV) 2010 as carried forward in JV 2020 form a vision for integrating doctrine, tactics, training, supporting activities, and technology into new operational capabilities. JV 2020 confirms the direction of the ongoing transformation of operational capabilities, and emphasizes the importance of further experimentation, exercises, analysis, and conceptual thought, especially in the areas of information operations, joint command and control, and multinational and interagency operations. Based on the joint vision implementation program, many capabilities will be operational well before 2020, while others will continue to be explored and developed through exercises and experimentation. The overarching focus of this vision is full spectrum dominance achieved through the interdependent application of dominant maneuver,

precision engagement, focused logistics, and full dimensional protection. Trust is the foundation of these four pillars – assisting the US Department of Defense (DoD) in maintaining national security. Maintaining national security requires the steady infusion of new technology, modernization and replacement of equipment, and the capacity for all in DoD to trust in network centric operations which will be dominated by speed and agility. (JV2020, 2000) However, material superiority in the kinetic solution alone is not sufficient. Of greater importance is the development of doctrine, organizations, training and education, leaders, and people that effectively take advantage of new technology that offers non-kinetic solutions.

As changes take place over time, it will likely be necessary to carefully examine the aspects of the human element of command and control. Leaders will need to analyze and understand the meaning of unit cohesion in the context of the small, widely dispersed units that are now envisioned. Decision makers at all levels will also need to understand the implications of new technologies that operate continuously in all conditions when human beings are incapable of the same endurance and be able to trust and focus their decisions based on the trust placed on information provided to them by both known and unknown sources.

As new information technologies, systems, and procedures make the same detailed information available at all levels of the chain of command, leaders will need to understand the implications of this on decision-making processes, the training of decision makers at all levels, and organizational patterns and procedures. There exists a potential

for over-centralization of control and the capacity for relatively junior leaders to make decisions with strategic impact are of particular importance.

As leaders in the military, or any organization for that matter, we can follow certain guidelines to develop the trust that is so critical for success in today's turbulent environment. Building trust is like developing physical fitness. We know we need to improve our muscle tone and "get in shape for the next semi-annual physical readiness training (PRT) exam" but find our good intentions easier to state than to carry out. We look for specific routines to fit into our busy schedules and then struggle to follow them consistently to get the results we want.

Developing fitness in relationships by building and maintaining trust is as simple-and as difficult-as getting in shape physically. The process involves exercising the muscles of power and vulnerability until both are equally strong and then using them in an authentic, consistent, ethical way to benefit our coworkers, our organizations, and ourselves. And trust, once developed through that kind of effort and persistence, will continue to grow and expand.

Trust in relationships is dynamic: it will continue to change as one's relationship changes. Trust is the glue of relationships. It breeds authenticity and respect. It encourages telling it like it is-communicating openly and honestly, fostering continuous improvement and growth. It allows errors to be made and corrected without undue risk of career destruction or personal embarrassment. The process of developing and maintaining trust is no different. Like a good exercise program, building trust involves some basic steps, hard work, commitment, perseverance, and continuing effort to achieve

the payoff. What is the "payoff" from trust? A division, department, command or even an organization, comprised of individuals connected by a bond of trust, is one that is flexible, adaptable, resilient, responsive to changing environmental demands and competitive maneuvers, and firm when necessary. One who has ever experienced a successful command or organization could easily agree with the above mentioned. So how do we develop trust-in others and ourselves? The "formula" for trust is easy to articulate but can be difficult to implement. Developing mutual trust means shifting our personal and organizational mental models from competition to collaboration giving up familiar, long-held assumptions and ways of thinking and doing. We may also need to reorganize the values, beliefs, axioms, and theories that together make up these mental models-models which have evolved throughout our lives and which form an essential part of our identity.

How, then, do we bring about such a shift in thinking? Most leaders start with a concept of self as being one who is in control, decisive, certain and directive--characteristics that followers have traditionally also demanded in a leader. Some leaders, however, realize they do not know all of the answers. They are caught in the "Catch-22" situation of needing both to display a public sense of clarity and control and to allow themselves to become learners, thereby trusting others just as they expect others to trust them. But for mutual trust to develop, people need to know whom to trust, when, and to what degree-and the choice must be made wisely. (Wyatt, 1996)

Developing mutual trust is no easy task. It calls for a leader to rely on his or her own expertise in one situation and on someone else's in another. In order to know when

to lead and when to allow someone else to lead, to have the insight to set boundaries that are neither too tight nor too loose to maximize creativity and minimize the degree of risk to the organization, are key attributes of good leaders. Creating vision, values, purpose, and mission while clarifying expectations reduces fear and builds trust. This can be seen in the Commanding Officer's Standing Orders or in an organization's mission statement.

What then is trust? Trust is the confidence or faith that people have in each other. It may seem that commitment and trust are only important or possible in a long-term relationship. For a short-term job assignment, people tend to feel they can overlook the relationship and use the pressure of the work to mask any deficiencies in commitment or trust. Trust can also be perceived to only happen when people know each other for a long time, that time is necessary to know each other's values and to rely upon the behavior of the other. It takes a special Leader to command people to have commitment and trust.

Coercion may get people to work together, but they will not form a partnership. A true partnership requires the free will of each participant to succeed individually and for the team.

There is no simple or single method for building trust. It is not merely a matter of working together for a long time. It does not occur if the partnership faces complexity and ambiguity and gives mutual support only to vent frustration about the shared adversity. It does occur through the attitude and intention of each partner. The leader can support the growth of trust by acting from a complete trust in the ability of each person. Such a situation occurs when each person feels that he or she still has the trust of the leader, despite mistakes or errors, and that the leader will provide personal support

even when there is a need to be critical of performance. Trust does not exist when judgment is present.

A good leader creates an environment for the partnership to develop its own path rather than acting as the leader and delegating tasks regardless of the overall complexity and ambiguity that the group faces. Trust is not present in a leader who delegates responsibility and then feels compelled to continually check on the work. (This is not to be confused with Management by Walking Around (MBWA).) One lacks trust when one worries or has sleepless nights thinking about the actions of the members of the partnership. Trust does not exist when fear is present.

We feel the leader's role in building trust is to exhibit complete honesty. Tell people what you are thinking, share information that you possess and they want to know, tell them when you cannot share information and why, ask and accept challenging questions, give clear responses or make a clear commitment to a future response.

E. IS TRUST TRUSTWORTHY

Network information systems (NIS) are more and more prevalent. Examples include: the public telephone system and the electric power grid. At the same time, these systems are neither trustworthy nor dependable. Moreover, the alarming trend seems to be for people, companies, and governments to be more and more dependent on NIS. Trust in a NIS can occur when the system does what it is intended in spite of the following: malicious attacks, environmental disruption, software bugs and/or erring operators. Contrary to what might be assumed, hostile attacks are today actually the least significant cause of system crashes and problems. It is important to keep this in mind.

Systems need to be able to tolerate squirrels and back-holes (chewing and slicing through wires respectively, as has happened) as well as attacks perpetrated by ill-intentioned hackers. Environmental disruption, operator-error, and insider attacks are the biggest sources of problems followed by buggy software, and only then, external attacks.

Trust in NIS building software to be trustworthy involves more than assessing functional requirements--what outputs must be produced for given inputs: one must consider nonfunctional properties of the system. We may not be told what attacks to expect so the specification of the problem is inherently incomplete. The timing of attacks can be unpredictable. Any attempt at representing attacks using predictive models could rule out possible attacks and, therefore, be incorrect. Consider the four trustworthiness dimensions discussed previously (i.e., malevolent attacks, software bugs, user error, and environmental disruption). These are intrinsically different from functional requirements; these are sometimes referred to as "negative properties." The challenge we face is that we have an open system, one in which some components are unspecified, and yet we are required to reason about all instances of the unspecified components; for example, we must reason about how the system would behave under a hostile attack without knowing what form this attack will take. Trustworthiness is a multi-dimensional problem. Is it possible that all of these four dimensions are really the same thing? After all, an environmental disruption can be seen as a random perturbation of the system and each of the other dimensions produces perturbations, so are they not all closely related? At a very coarse level they can be seen that way, but closer study reveals that they are different.

Environmental disruptions are events that are uncorrelated. If events are independent, then it may make sense to use replication in order to build a system that will tolerate some number and frequency of failures. Hostile attacks, on the other hand, are correlated. Replication does not work for correlated failures. Operator error is in some ways even worse than a hostile attack, because operators are often trusted users who will have privileges that outside attackers will not. Moreover, software bugs are worse than operator error because the buggy software may have arbitrarily high levels of privilege.

It may be useful to know why NISs are becoming more prevalent and what is driving that process. In the private sector, organizations today seem to be driven by the need to operate faster and more efficiently. Profit margins are thinner and expectations are high. For example, consider just-in-time manufacturing wherein inventory and material are not warehoused but instead shipped to arrive exactly when needed. In this kind of environment, timely information (e.g., who needs what and when?) becomes essential, thus the need for network information systems.

In the quasi-public sector there is a new climate of deregulation. Less regulation produces competition, which produces a need for increasing levels of productivity. Companies thus need to lower expenses, and one way to do this is to decrease excess capacity (e.g., power reserves, bandwidth). Lower excess capacity results in the need for finer control over the existing capacity, which in turn requires a good supporting information systems. Lower excess capacity can result in less trustworthiness by creating a less stable system. Excess capacity can, in some cases, take up the slack in the event of a system failure or disturbance. With less "slack" it becomes more likely that a "small"

failure could have large repercussions. Another result of the need to lower expenses and attract customers in deregulated industries is the introduction of new and complicated features to existing services (e.g., in the telephone industry consider things like call-forwarding and caller identification). The more complicated a system becomes, the less reliable it will be. The addition of features increases complexity, which may well result in unanticipated and undesirable behavior. In telephony, this is known as the "feature interaction problem."

The development of new industries exploiting NIS, such as electronic commerce, is a third reason for the growing prevalence of such systems. In short, it seems as if we are heading towards a situation in which there will be many untrustworthy network information systems. This problem will need to be fixed, but one might ask: How bad is this problem? What are its dimensions? The consequences of untrustworthiness include denial of service (DoS), which we have recently been seen in Silicon Valley with attacks on internet-based business Yahoo and Ebay. Yet availability can be extremely important. Telephone and power outages can result in loss of life and civil unrest. Information disclosure is another problem. It can result in personal embarrassment, financial loss, or even loss of national security. Information alteration is yet another possible problem that can obviously affect everything from a student's grades to the nation's economic health.

All of this adds up to a relatively new form of warfare termed "information warfare." Such warfare can be overt or subtle, ranging from interfering with military communication to planting sleeper programs to manipulate the stock market. Information warfare opportunities exist only because we as a society are so dependent on information

systems. It is possible to attack anonymously without ever being physically present. Why don't trustworthy systems exist? The next obvious question is, given all of the above, why are network information systems not built to be trustworthy? One answer is that it is not clear in all instances how to make systems trustworthy. However, it is likely due to many factors, including direct and indirect costs.

F. TRUST IN THE INTERNET ENVIRONMENT

Historically, the concept of trust has always been important, but it was usually defined in a subjective manner, such as a feeling of comfort. Philosophy and religion expanded the notions of trust, but rarely was trust defined within a quantitative or qualitative framework. The Internet is inherently an untrustworthy medium due to the fact that a user is often uncertain with whom he or she is communicating; the same can be said for computer processes of intelligent agents, which act as proxies for the user. Additionally, the user has very little control over the software and hardware programs that are executing on his/her behalf.

Internet certification protocols attempt to deal with the concept of trust without ever defining what trust is. Without a formal and commonly accepted definition and identifications of the components of trust, how can a protocol effectively or efficiently address the issue of trust? In order to incorporate trust into electronic commerce, public key cryptography, and basic communication, one must understand and effectively manage trust.

Trust can be thought of in terms of faith or confidence. If a ladder looks wobbly, one is unlikely to trust it to hold one's weight as one climbs up to the top rung of the

ladder. Now consider trusting the security of the Internet. If the Internet mechanisms for enforcing authentication, authorization, privacy, integrity, and non-repudiation policy do not appear sturdy to the users, then users will hesitate to use the mechanisms. Trust can be lacking for reasons both real and perceived. One of the reasons there is not a high level of trust in the Internet for conducting financial transactions is that people simply do not understand the enabling technology.

The confidence in the ability to authenticate companies and individuals over networks is critical to commerce and defense. If you did not trust that when you walk into your bank, you were truly dealing with authorized agents, banking would take considerably more time than it does today. It is easier to authenticate companies and individuals with whom you have an established relationship or who are introduced to you by someone you trust.

What happens in regard to new relationships? Generally, these are authenticated by some third party. You must trust the third party or there is no confidence in the authentication. Trust must also be present for privacy measures. People will not bank with a financial institution that does not employ adequate safeguards to protect their client information, such as signature, date of birth, social security number or tax identification number, and financial information (e.g., total debt or savings). The same is true of data integrity. If a consumer fears that a hacker might alter their credit history, the consumer might not apply for credit.

In the network world, the requirements for trust are as yet not well defined or agreed. Trust is something that has yet to be firmly established for network usage. The

lack of confidence is part of what is behind the concern that Internet users express as a need for security. When one places a high-value transaction on the Internet, a user will have established some form of trust in advance prior to entering one's credit card number. One may not have any reservations when entering a credit card number to purchase a book from a reputable company like Amazon.Com. On the contrary, one would typically not be so inclined to submit a credit card order to a company that the user has no previous knowledge about.

We must not only have a secure Internet, we must convince consumers and businesses alike that the level of security is trustworthy. This will require education, experience, and infrastructure. Users need to understand how Internet security works and the safeguards that protect them.

Trust also deals with assumptions about expectations and behaviors. This implies that some aspects of trust cannot be measured quantitatively, that there is a risk associated with trust, and that the determination of the level of trust to be placed in someone or something cannot always be fully automated. However the concept of a trust model is useful because it shows where and how trust is initiated in a PKI, which can allow more detailed reasoning about the security of the underlying architecture as well as intentions imposed by architecture. The term 'trust' is frequently used in another way that is useful to us as well. The literature about PKI often refers to so-called *trusted public keys*. This phrase does not describe assumptions and expectations about behavior. Rather, the public key is said to be trusted by a user when the user is convinced that the

public key corresponds to the private key that legitimately and validly belongs to a specific named entity.

There are a number of trust models that represent attempts to define trust, assign values, choose alternatives, evaluate alternatives by measuring the attributes, and select the best alternative. By developing a new paradigm for expressing trust in a PKI setting, our trust model will distinguish itself from its peers so it can be effectively used with automated security protocols. In this thesis we will evaluate the various trust models and select those concepts of a trust model that are appropriate for any organization, and then through the use of a case study, discuss the concepts in the context of a network centric warfare information infrastructure.

G. SECURITY CLEARANCE

The task of protecting sensitive information has always been informally described as one of uncertainty and risk. This is true whether the information is held by people or by computers. If one is allowed to handle Secret information, for instance, one can in general communicate that information to anyone else cleared for Secret and access any computer system whose accreditation range includes Secret. Even though a given computer system is allowed to handle Secret information, it cannot be allowed to communicate that information to other computer systems, even if they are allowed to handle it. This is an example of “need-to-know” policy where we do not trust everyone who has a Secret clearance to see all of the objects labeled Secret. For example, I can see information related to a project I am working on, but I may not be trusted to see compartmented information about another project which I don not work on.

The clearance process then is modeled as one of restricting the population to which the probability distribution applies. A person is to be granted a clearance of a given level, if the expected distribution of damage, computed over all people granted that clearance, during their tenure, assuming they are granted access to information of that level, falls below some designated threshold. It is further noted that a person holding classified information can cause damage in two ways: by directly acting on the information in some way, or by passing it on to someone else.

Whenever any sensitive information is given to anyone, there is a risk that it can be misused and "cause harm or damage." (If it could be made accessible to anyone without any reservation it would not be sensitive.) "More sensitive" information, for example, of higher classification, is more sensitive than "less sensitive" information, for example, of lower classification. This is because the risk of damage to the national interest being done by improper use of the more sensitive information is greater than that of it being done by improper use of the lower.

The more sensitive (more highly classified) information is, the more confidence there must be that the person it is given to can be trusted to deal with it. The reason I do not give TOP SECRET information to someone, only cleared to SECRET, is not that I know that he or she can not be trusted with it, but that, in effect, I have been told not to trust my ability to interpret DoD security policies. The security authorities could also decide with sufficient confidence that he or she can be trusted, since the clearance process at the SECRET level, and all the other doctrine going along with it, is not thorough enough. A person is allowed to have access to a given piece of sensitive

information because, in a very subjective sense, the risk (danger of improper use) inherent in the sensitivity of the information is balanced by the risk in the personnel security process. The more risk there is in dealing with a particular kind of information (i.e., the higher its classification), the less risk there must be (higher clearance) that the decision that someone can be trusted with it is in error.

Since the typical possessor of classified information cannot be expected to have the background to perform the above act of balancing risks, the security authorities do so for him or her ahead of time by assigning classification and clearance labels. A clearance label (e.g., SECRET) is assigned to designate each particular clearance process. Information is given a classification label such that if when information of that label is given to someone with the same or higher classification the risks are adequately balanced, but not if it is given to someone with a lower classification level.

H. TRUST IN DIGITAL SIGNATURE

If one was to look back, businesses and individuals used letters of introduction to vouch for their identity. Our driver's license, military identification, or passport also substantiates that we are who we portray to be. More recently, the growing use of electronic transactions has led to the use of a digital signature to authenticate the identities of parties to transactions. Or do they? The digital signature is effective when a user has been issued a certificate of identity that has been digitally signed by a trustworthy certificate authority. And there's the rub, because we have not come to terms with the question, which should serve as a certificate authority? Or, for that matter, who should certify the certificate authorities?

The basic technology questions concerning digital signatures and certificate authority software have been settled; software to issue certificates safely and to track expirations or revocations is available from several vendors. Some base their designs on a hierarchical chain of CAs (a root authority is authorized to certify other CAs, who in turn are authorized to certify the identify of specific end users or user accounts). Others base their design on a "shared trust" model (Gogan, 1999), in which participants who have been issued certificates of identity can, in turn, vouch for the identities of parties known to them. From a technical standpoint, either approach will work, with some compromises. For example, some experts say that once there are lots of participants in a hierarchical chain, the throughput times will be unacceptable unless a design breakthrough is achieved.

Others say the shared-trust model comes with a greater security challenge, since it is far more difficult to revoke a certificate of identity that has been given to an impostor by a party who thought they knew the person's true identity. The tough questions revolve around management and policy implications, not technology. A business might choose to serve as its own certificate authority, using software from Microsoft, Netscape, Baltimore Technologies, Xcert International, Pretty Good Privacy, and others. Alternatively, a firm might hire Equifax, VeriSign, or GTE to operate a CA service for them. (Gogan, 1999)

II. BACKGROUND

A. TRUST

There are many ways of describing trust. Audun Jøsang defines trust management in his paper, *Trust Management for E-Commerce*, as follows:

"In the context of e-commerce and IT security we will define trust in principals as the expectation or belief that they will behave according to a given policy and without malicious intent, and trust in systems as the expectation or belief that they are secure and will resist malicious attack. Trust is thus a belief and we assume trust to be based on evidence, experience and perception." (Jøsang, 2000)

Jøsang refers to trust in the physical world as trust in things and in other people that is based on our experience with them, information we have received about them, and how the people appear to us. He mentions that trust is a very subjective phenomenon, meaning that one does not necessarily trust the same things or the same people as you and vice versa. The number of people we can potentially relate to from face-to-face contact is also limited by distance and physical constraints. On the Internet, on the other hand the number of people that are on-line only limits the number of people we can potentially come in contact with. (Jøsang, 2000)

In his doctoral research project, Dr. Jøsang defines trust as follows:

"From an information security point of view, human agents are trusted because they are believed to be honest whereas systems or entities are trusted because they are believed to be secure. Trusting a human agent is then simply to believe that the agent will cooperate during an interaction, whereas trusting a system entity is to believe that the entity is resistant to malicious manipulation. Trust must therefore be seen as a belief, and an important part of the work has been to develop a new belief model and related calculus called subjective logic, as none of the existing belief models were found suitable." (Jøsang, 1998)

When conceptualized as a psychological state, trust has been defined in terms of several interrelated cognitive processes and orientations. First and foremost, trust entails a state of perceived vulnerability or risk that is derived from individuals' uncertainty regarding the motives, intentions, and prospective actions of others on whom they depend. (Kramer, 1999)

One noteworthy paper on trust is called "The Meaning of Trust" by Harrison McKnight and Norman Chervany. These two authors have noted that the meaning of conceptual trust must be resolved. Their paper acknowledges that trust lacks the consensus of what trust is, although many scholars provide definitions supporting their theory on trust. McKnight and Chervany have reviewed and tabulated types of trust concepts from research articles and books. Their results show a bewildering array of meanings and connotations for trust. Of the sixty articles, thirty of them state that trust refers to a perceived attribute, or set of attributes, of the person trusted. And of those cited, the sources generally define trust in terms of beliefs or expectations about the other person. (McKnight and Chervany, 1996)

Their research indicates several things about the trust literature. First their research results indicate that trust is often defined in terms of expectancies or beliefs. Beliefs reflect perceptions about the role the other party plays in trust relationships. Second, many definitions include affective, or cognitive/affective, aspects. These definitions of trust typically include a phrase about feelings of security about, or confidence in, the trusted party. Third, it is noted, that a large number of definitions refer to trust as a behavior. (McKnight and Chervany, 1996)

Also noted in their table of references, is the considerable breadth of coverage of the types of definitions. If this were a test of consensus on trust definitions, then researchers would receive a very low consensus rating. They found that thirty-six of the sixty articles or books (sixty percent) define trust in more than one conceptual category. On average, these researchers used 1.9 categories. Hence, most of these individual researchers feel trust has more than one meaning. This is additional evidence of the breadth of the trust concept. (McKnight and Chervany, 1996)

One advantage of conceptualizing trust in terms of choice is that decisions are observable behaviors. Another is that organizational theorists possess a well-developed conceptual idea for pursuing the theoretical and empirical implications of trust-as-choice.

Hardin's conception of encapsulated trust captures some of the essential features of this view. A rational account of trust, he notes, includes two central elements. The first is the knowledge that enables a person to trust another. The second is the incentives of the person who is trusted (the trustee) to honor or fulfill that trust. Individuals can trust someone, Hardin proposes, if they have adequate grounds for believing it will be in that person's interest to be trustworthy "in the relevant way at the relevant time" (Hardin, p. 153). This notion of trust, he observes, is not predicated on the individuals' narrow contemplation of their own interests but is enfolded instead in a sophisticated understanding of the other party's interests. "You can more confidently trust me," (Hardin, 1991) posits, "if you know that my own interest will induce me to live up to your expectations. Your trust then encapsulates my interests" (Hardin, p. 189).

Intuitively, a minimal level of trust would appear to be necessary for any negotiated transaction to occur. However, what exactly is meant by "trust?" every scholar uses the definitions that we will offer in this paper.

A related, and perhaps more fundamental question is, "Why is trust important?" We believe that trust is worthy of consideration because if we are vulnerable to one and another or are considering an option that makes us vulnerable to one and another, then if we can trust the other, we do not need to worry about exploitation by the other. Otherwise, we must protect ourselves from the other or avoid decision options that make us vulnerable to the other. The issue is whether we believe the other will act in our best interest. Although there is much in common among various conceptualizations of trust, trust is a multi-layered phenomenon, and there are differences in the approaches taken by various theorists. Some view trust as a dispositional variable. Others view it as a temporary state, and relate it to various situational antecedents (e.g., cooperation) and consequences (e.g., integrative bargaining). Some perspectives emphasize that cooperation is a sufficient operationalization of trust. Others emphasize the development of trust through a series of predictable, cooperative behaviors. Still others stress the transitory psychological state of trust. It is hoped that this paper can clarify some potential confusion about this construct and that in the future, researchers can use some common terminology. Finally, we attempt to place the construct of trust into a comprehensive framework as it relates to mandatory and discretionary controls, building upon and incorporating previous models of trust, which have often dealt with only selected portions of this process.

One perspective treats trust as a stable dispositional variable or individual difference variable. Thus, some persons are "more trusting" of others. This perspective treats trust as an individual difference construct of "general trust" toward other people and assumes that the negotiator (subject) will usually react similarly (trustfully or distrustfully) toward all opponents.

B. THE TOFFLERS AND THE WAVE OF TRUST

Our thesis would not be complete if it did not relate trust to the Tofflers who have influenced those of us who study Information Operations. The United States and many countries throughout the world, when talking about society, technology, the future, and future trends, one cannot help bring out the nonfiction writings of Alvin Toffler. One of his books, *The Third Wave*, addresses questions of why so many changes are occurring so rapidly and what those trends may mean for the future. (Toffler, 1991)

One who does not look ahead at the trends in society, technology, business, global competition, criminal justice systems, crime, and any associated rapid changes will likely have a stagnant information security program that fails to meet the needs of the business or government agency.

Is there a relationship between what is happening in the technological arena and its accompanying rapid social and business changes and the need for trust and more security in the information world? There at least appears to be a relationship between them. This technology is changing rapidly, twenty-four hours a day.

Toffler speaks of societies of a world going through or about to go through three "waves." The first wave is the agricultural revolution, which has taken thousands of

years to develop. According to some legal experts, this period, at least in the United States, started with the beginning of the human race to about 1745. Obviously, agriculture is necessary for humans to survive. During this period, people live in small and sometimes migratory group, feeding themselves from fishing, foraging, hunting, and herding. Subsequently, each migrated into clusters, and towns, then cities.

During the first wave, information was passed by word of mouth or in written correspondence, which was usually sent by courier. People were more dispersed and transportation was primitive. This meant that there was less communication among people. During this period, the number of people who could read or write was relatively few in comparison to the total world population. Therefore, protection and trust in information, albeit very human intensive, was not a major consideration as it is today.

Such threats, as theft of information in the written form, were minimal, because most people of the world could not read or their ability to read was very limited - although they could destroy the written message. Perhaps, this type of destruction was the first instance of denial of service. Information verbally relayed could be misinterpreted or changed, a method that still poses a threat to information security for today's IT society.

Information security in those days was much less difficult in comparison to today's standards. In those days, a king who did not trust someone and was afraid he or she was going to disclose sensitive information to other people, cut out the person's tongue. As people became more educated, learning to read and write, trust and information security challenges broadened. (Toffler, 1991)

The second wave, Toffler calls the "rise of industrialization," took less than 300 years. This was the age of steel mills, oil refineries, textile plants, mass assembly lines, and the like. The people came together to work in these industries. This period lasted until just a few years after World War II. In the United States, its decline, according to Toffler, is believed to have started around 1955 when, for the first time, white-collar workers outnumbered blue-collar workers.

The second wave involved the building of the great cities of the world, the period of great inventions like the telegraph, telephone, air transportation, and computers. This introduced an increase in education, mass transportation, the exponential growth in communication: the sharing of information. (Toffler, 1991)

The dissemination of information improved and increased with the invention of communications systems and increased consolidation of people into large cities. This also made it easier to educate the people needed to work in more modern factories and offices of the period.

Sharing information through various communications channels introduced new challenges for protecting information. The primary information security protection methodology that came into being during this period was cryptography. Cryptography was applied mainly by government, to protect information transmitted electronically. Although businesses were beginning to look at the use of computers, most computers were cost prohibitive and these systems were operating primarily in a stand-alone mode. In other words, the computer did not talk to other computers.

For much of this period, information security for business and government agencies consisted of personal security and some minimal level of physical security. As the computer became more sophisticated, the main protection mechanism used for computers changed very little. After all, why worry about such things as access control other than physical security? Not many people knew how to use computers in the first place. At the beginning of this time, very few people worked in the computer field, so the human threats to information systems and their information was relatively manageable compared to today. Therefore, at first, the threats to information systems and their information were small.

The third wave, the age of technology, information, and knowledge, is sweeping across the earth and will have done so in decades not centuries. This period, which we are in now, has produced more advances than the first and second waves combined. This period has seen the rapid growth of technology that is playing a major role in our changing world. (Toffler, 1991)

Today, because of availability, power, and low cost of microprocessors, society is building the global information infrastructure (GII). GII is the massive international connection of world computers that carries business and personal communication as well as that of social and governmental sectors of nation states. It has connected entire cultures, erased international borders, incubated "cyber-economies," established new markets, and changed our entire concept of international relations.

C. GENERAL SECURITY ISSUES

In the real world, security decisions are based on three things:

- value,
- locks, and
- police

We try to buy good enough locks so that the "bad guys" cannot break in too often. The terms "good enough," "break in," and "too often" are key. We also assume that the police and courts work, so "bad guys" are caught and punished. "Police" in this context is a generic term for any person or group that might pursue offenders; it includes the corporate hierarchy and the legal system. Similarly, the "bad guys" could be anyone, anywhere, including system operators for the system being secured. By "often enough" we do not mean *always* but enough so that crime does not pay. In other words, the expected gain from committing a crime must be negative. Value is an important aspect of this characterization, because generally we do not protect things of little value.

A constraint we place on any security mechanism is that it adds a minimum amount of interference to daily life. For example, locks must not be difficult or annoying to use. If they are, it's likely that people will find ways to circumvent the annoyance, and thus nullify the security protections the locks offer. It should also be noted that, with rare exception, is a security breach a catastrophic event. Risk management supports planning for recovery from a security breach and decreases the need for complex and annoying locks. (i.e., locks that pose a hindrance to conducting business) For example, rather than installing a complicated locking system for automobiles, we buy auto insurance to help deal with costs that arise in the event of damage or theft.

Externalities also have a role to play. Briefly, an externality occurs when somebody or some agency does something in which the cost implications for the doer are not the same as (usually significantly less than) the cost implications for society. For example, think of companies that pollute the environment. The cost of cleaning pollution is usually great, and until recently there was no corporate penalty for *not* fixing a pollution problem. In short, an externality exists when it is cheaper to do the wrong thing. This has obvious large implications for security--an insecure subsystem may enable a system-wide attack of great consequence.

There are a number of things to observe. First, note that all locks are not the same. They typically have different keys as well as different strengths. The strength of the lock tends to be chosen according to the value of what is being protected. The environment also influences the type and strength of the locks being used as well. For example, apartments in well-known, safe neighborhoods are likely to have fewer and weaker locks than apartments where crime is customary. Second, people pay for security they believe they need. Security is not monolithic and there is not one mechanism for everyone. Security is scaled with respect to both the value of the thing being secured and the threat against it. People's security "needs" are usually based on the perception of what's going on around them. If your neighbor's home is being broken into, then it is likely that you will buy more security equipment than otherwise. Third, the police are central to the picture. The system still works even when locks are completely removed.

Locks are only a deterrent; however, it is essential that there be enforcement and punishment strategies in place. There will undoubtedly be some security breaches no

matter how good the locks are. Thus, it is critical that bad guys be found. Locks reduce temptation as well as reducing the police workload. Finally, security, as we have portrayed it, is holistic. It is only as good as its weakest link. Attackers will look for the weakest link, and thus it is generally best to expend effort in determining where the weaknesses are and shoring them up. Given limited resources, the best approach is to make all elements equally strong, thus eliminating weakest links.

D. APPLICATIONS TO COMPUTER SECURITY

We now move from an abstract discussion of security in our day-to-day lives to the world of computer security. How can the above discussion be applied in terms of computer security? With regard to computer security, the story is told in the following terms:

- *Vulnerability*: A weakness that can be exploited to cause damage.
- *Attack*: A method of exploiting a vulnerability.
- *Threat*: A motivated, capable adversary that mounts attacks.

Bugs in a software system are vulnerabilities. Since we are not really good at building large systems, it seems clear that any large software system will have many vulnerabilities. While a first strategy for addressing a security problem might be to find and fix each vulnerability, in fact, this is likely to be too costly to be practical. Rather, it is better to first identify *threats*, and then work on eliminating only those vulnerabilities that those threats would exploit.

As an example, consider the problem of intercepting cellular phone transmissions. This possibility is clearly a result of design vulnerability--a consequence of the way

cellular phone signals are encoded and transmitted. A threat that exploits this vulnerability would be the small number of people who want to do this and have the knowledge and equipment to intercept transmissions. When cell phones were first introduced, the equipment was hard to come by and few people had the knowledge to mount an attack. Thus, the threat was small. Currently, just about anyone can buy the equipment; the threat is huge. The vulnerability has remained the same, but the nature of the threat has changed. So, what about commercial products we can purchase straight off the shelf to protect us, especially if we are talking in terms of trust.

E. TRUST IN COMMERCIAL-OFF-THE-SHELF (COTS) PRODUCTS

COTS (commercial-off-the-shelf) products today dominate the software and systems markets. There is a huge economy of scale involved in building and using COTS components. Imagine someone in charge of integrating a large system, which needs to be completed on-time and on-budget. It is faster and cheaper to use COTS components, and this also can reduce some aspects of project risk. Another benefit of using COTS products is interoperability. Upgrading from one version of software to the next is usually straightforward and the easiest thing for users to do, even though there may very well be a better product available. As an example, the government, even the NSA, uses COTS equipment for all but its most secure communications. Those who provide COTS products (such as Verisign, and Entrust) know that the market prefers features over trustworthiness. However, this is changing.

It may be that the market and individual consumers are not really conscious of this, but it seems to hold just the same. It is generally not clear to consumers what

trustworthiness would provide, and the market is not aware of the risks involved. One counter-example is in the area of hardware failure. Fault tolerance is much appreciated by the market, perhaps due to the fact that failure of a machine has obvious and immediately impacts on productivity. On the other hand, in the past, there have been instances where the market did not necessarily appreciate something, and yet the manufacturers still provided it. Why do COTS producers not provide trust and trustworthiness?

They especially do not tell us whom we should and should not trust! The COTS market rule of thumb is that the earliest entrant to a market is the most likely to succeed. In other words, time to market dictates success. Implementing trustworthiness increases the time to market. It requires extra functionality, fault tolerance, better debugging, ways to provide assurances, and so on, which all add to development time and cost. In short, there is every incentive for COTS producers not to provide trustworthiness, and given the current climate of deregulation, it is not likely that the government will legislate requirements on trustworthiness any time soon.

Another reason for the lack of trustworthy NIS's is the existing communication infrastructure. Ultimately, the telephone companies have control, and they still function under a very old tariff system. This system does not encourage them to provide things like path-disjoint (i.e., more fault-tolerant) service. The internet today is very easy to "crash" with denial of service attacks. U.S. Government policy also does not encourage the production of trustworthy products, particularly with its restrictions on the export of cryptographic equipment and its push for key-escrow. Key-escrow was designed for the

American telephone network in the early 1990s to incorporate each telephone with a chip, known as the Clipper Chip, containing a secret cryptographic algorithm, also known as skip-jack algorithm. (Lubbe, 1998, Menezes, 1996) As a result of the lack of public trust, no incentive was given to manufacturers to build equipment which had strong cryptography and that could be trusted by everyone. This lack of incentive to manufacturers was a result of the public's lack of trust in the escrow agent, the escrow scheme, and in the strength of the cryptography. As a result, the U.S. Government backed off from mandating the use of the chip. All of the above paints a depressing picture. Are there any glimmers of hope to be found? The existence and prevalence of COTS products implies that if trustworthiness were to be implemented, then the prevalence of COTS would enable widespread deployment.

III. MODELS

A. FORMAL MODELS

This section builds on the discussion of models provided by Gaines (2000).

While it is appealing to approach models that are based on security, it is important to recognize that many good ideas turn out to be not so good on close inspection. The formal model is typically a subject-point-of-view model. The vision of our model is to create a technology for putting a rich set of policies in a system that incorporates both discretionary and mandatory trust. Our model, to be discussed in Chapter 4, will make it easier to state, formalize, and analyze discretionary and mandatory trust in order to increase the availability of diverse information infrastructures.

We will first look at the basic concepts that turn up in models, and particular, models of security. One must not interchange security with trust. They are different. As stated in Chapter I, trust is a cognitive function that is unpredictable. Trust in humans is typically not determined in terms of how secure they are. Although humans can be assigned a security clearance (i.e., the Commanding Officer is cleared for top secret), the trust placed in that individual to command was based on the person's characteristics that allowed him or her to attain that level of security clearance, not that individual's resistance to coercion to divulge sensitive information. As mentioned, security in humans is associated with an assigned security clearance. Security in an organization or systems, are associated with how well they are protected against "malicious attacks." (Jøsang, 1999)

We will then look at key distinctions among various kinds of security models, and finally how these concepts and distinctions are reflected in well known examples, and apply them to our discretionary-mandatory trust model.

B. BASIC CONCEPTS TO BE MODELED

A secure computing system may decompose into data structures, processes, information about users, I/O devices, and security attributes for controlled entities. The primary aspects of models include policy objective, locus of policy enforcement, strength of policy enforcement, granularity of user designations, and the locus of administrative authority. These policy aspects are normally reflected in security attributes of controlled entities when the policies are formalized. Security attributes may be implicit. Meaning, they need not be directly implemented in the data structure.

1. Strict Hierarchy

One model which appears to have numerous names, is referenced to as a "strict hierarchy" of Certification Authorities (CAs). It is an inverted tree with the root at the top of the tree, and branches extend downward, with leaves at the bottom (Adams, 1999). Figure 1 shows a strict hierarchy of CAs trust. In this rather simple model of an inverted tree, the root represents a particular CA, commonly known as the root CA.

The root is known as a "trust anchor" for the entire domain of PKI entities which spread out to other entities below. Below the root are different levels of CAs. These CAs are also known as subordinate CAs because they are subordinate to the root. (Adams, 1999) These are represented by the intermediate nodes in Figure 1, from which

further branches can branch out. The bottom of the tree is also known as non-CA PKI entities, which Adams refers to as *end-entities* or simply *end-users*.

The term "root" while creating a mental image of the beginning of a tree with branches and leaves actually portrays something more fundamental. It is not simply a starting point for a network, communications, or subordination architecture; it is a starting point for what we shall call "trust."

All entities in this tree community hold a public key as their trust anchor, which is the starting and ending point of their trust anchor, their starting or point of trust for all certificate-verification decisions.

In this model, all entities in this hierarchical tree, must trust a single entity, or root CA. The root CA certifies (that is, creates and signs the certificates) for zero or more CAs immediately below it. Each of those CAs certifies zero or more CAs immediately down from it. At the second to last level, the CAs certify end-entities.

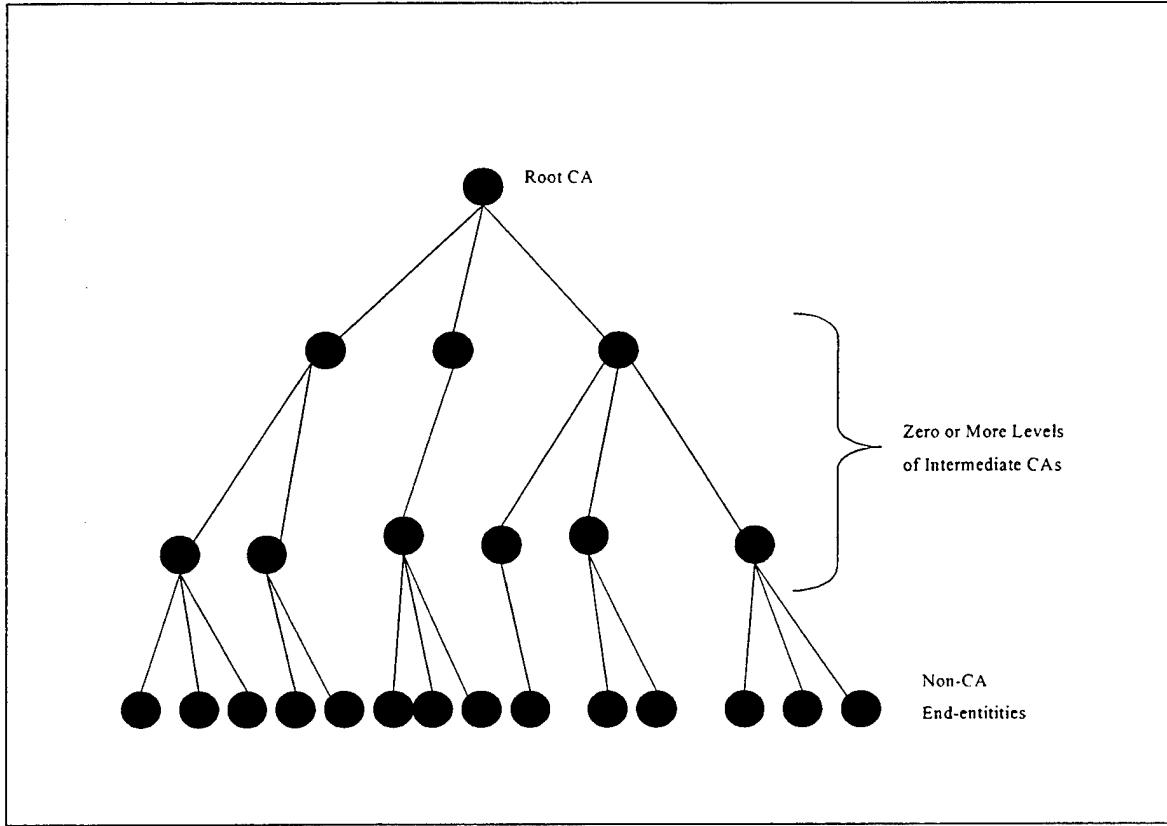


Figure 1, Strict Hierarchy (from Adams, 1999)

In a strict hierarchy, all entities hold a trusted copy of the root CA key as an anchor, that is, a starting point or ending point for certificate path processing. Each entity in a strict hierarchy (both intermediate CA and non-CA leaf) must be supplied with a copy of the root CAs public key. This public key installation process is the foundation to the certificate processing for all subsequent communication in this model. Therefore, one may acquire this key through physical channels such as paper, mail, telephone, or simply, electronically through the internet. The key must be confirmed, usually through some form of a hash, which may be sent again by courier, mail, telephone, or the internet.

Also note, that in this strict hierarchy environment, end entities are certified, that is issued certificates by, the CA immediately above them. A strict hierarchy is not appropriate for every environment. For example, this trust model will not work in most civilian organizations due to the lack of structure and discipline. Contrary to the civilian organization, some environments exist solely in a hierarchical form. An example is one of both theory and practice established in the military.

To establish secure communications between two end entities, Alice and Bob, a number of specific events must happen. First, Alice, holding a trusted copy of the root CA public key, must verify the certificate of another entity, Bob. Bob's certificate is signed by CA₂, whose certificate is signed by CA₁, whose certificate is signed by the root CA. Alice, with the root's public key K_R, can verify the certificate of CA₁ and, therefore, extract a trusted copy of CA₁'s public key K₁. Then this key can be used to verify the certificate of CA₂, which leads to a trusted copy of CA₂'s public key K₂. The key K₂ can be used to verify Bob's certificate, leading to a trusted copy of Bob's public key K_B. Alice can now use the desired key K_B, depending on its type, to encrypt messages for Bob, or to verify digital signatures Bob created (Adams, 1999).

2. Distributed Trust Architecture

Contrary to strict hierarchy, in which all of the users in the PKI community trust a single root CA, the distributed trust architecture, figure 2, distributes trust among two or more CAs.

As an example, Alice may hold a copy of a public key of a CA, as her trust anchor, and Bob may hold a copy of a public key CA, as his trust anchor. Because these

CA keys serve as trust anchors, it follows that each corresponding CA is a root CA for a strict hierarchy involving some subset of the total PKI community (CA₁ is the root hierarchy that includes Alice, and CA₂ is the root hierarchy that includes Bob.)

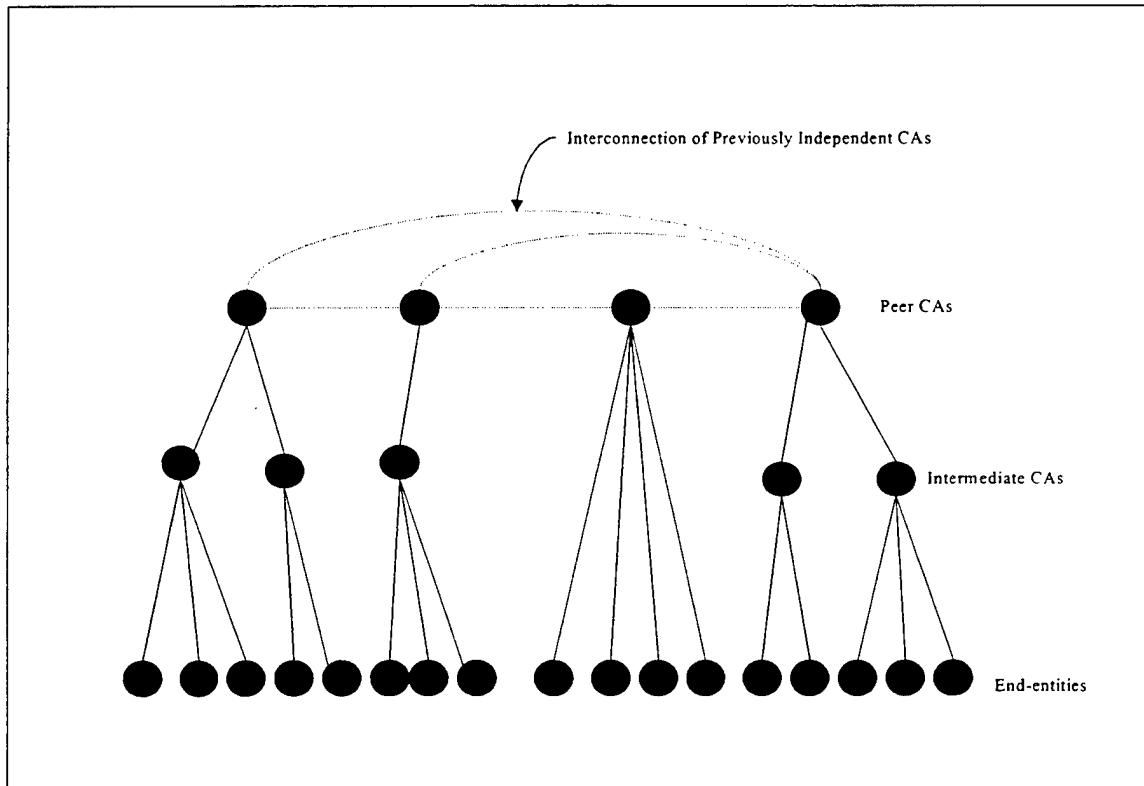


Figure 2, Distributed Trust Architecture, (from Adams, 1999)

This distributed trust architecture may be referred to as a fully-peered architecture, because all the peers are independent peers. Root CAs are peers with each other, but the root CAs acts as a superior for one or more subordinate CAs. An example of use of this architecture is multiple interconnecting independent, pre-existing PKIs from different organizational domains. This architecture supports the establishment of secure

communications through existing certificates. This form of passing certificates is also known as cross certification or PKI networking.

In today's electronic age, many enterprise domains deploy their own PKIs, and PKIs do not necessarily emanate from a common root CA. PKI domains can be configured in a number of ways. The process of interconnecting the peer root CAs is commonly known as cross certification. Some may know this as PKI networking which is growing in use. While on the subject, we must discuss several configuration forms of employed in cross certification. Two different kinds of cross certification are mesh and hub and spoke.

a. Mesh Configuration (Distributed Trust Architecture)

A mesh configuration is also known as "Distributed Trust Architecture" (Adams, 1999). In this architecture, all root CAs are potentially cross-certified with each other. In other words, two root CAs will cross certify whenever their respective communities need to communicate securely.

Figure 2 represents a partial mesh hybrid distributed trust architecture. It is not a full mesh because no direct cross certification agreement is in place between the first and third CAs.

b. Hub-and-Spoke Configuration

In the hub and spoke configuration, each root CA cross certifies with a single central CA whose job is to facilitate such interconnection. Adams refers to this central CA as a hub CA with spokes that branch out to the various other root CAs.

Adams calls is also known as a bridge CA - bridging communication gaps between pairs of roots.

The attraction of this configuration is that when fully connected, the configuration of each fully connected entity requires the same amount of cross certification agreements for the root CA. This is because each root CA cross certifies only with the hub (Adams, 1999).

The hub CA should not be viewed as a root for all the systems that cross-certify with it. The hub and spoke configuration does not create a hierarchy. The fundamental difference between these two trust models lies in which keys end-entities hold. In the hub and spoke configuration, no entity holds a hub CA key as an anchor. Instead, each entity holds a trusted copy of the key of a CA in its own domain and, through certificate path processing, obtains the key of the hub CA, and then a CA in another domain, and eventually the key of the target end-entity in that domain (Adams, 1999).

3. Web Model

The web model, figure 3, comes from the World Wide Web and the dependence of popular Web browsers such as Netscape Navigator and Microsoft Internet Explorer. A number of CA public keys are pre-installed in a standard, off-the-shelf browser. These keys define the set of CAs the browser user will initially "trust" to act as roots for certificate verification.

It is generally recognized that few users know enough, with respect to PKI and security issues, to change, understand, or modify this aspect of browser behavior.

By interconnecting with relevant domains, the Web model instantaneously makes Alice, for example, a relying party of all domains represented by the browser.

Each browser vendor has its own root, and it certifies the "root" CAs that are embedded in the browser. The only real difference is that the root CAs, rather than being certified by the browser vendor's root, are physically embedded in software releases as a means of attaining secure bindings between CA names and their key.

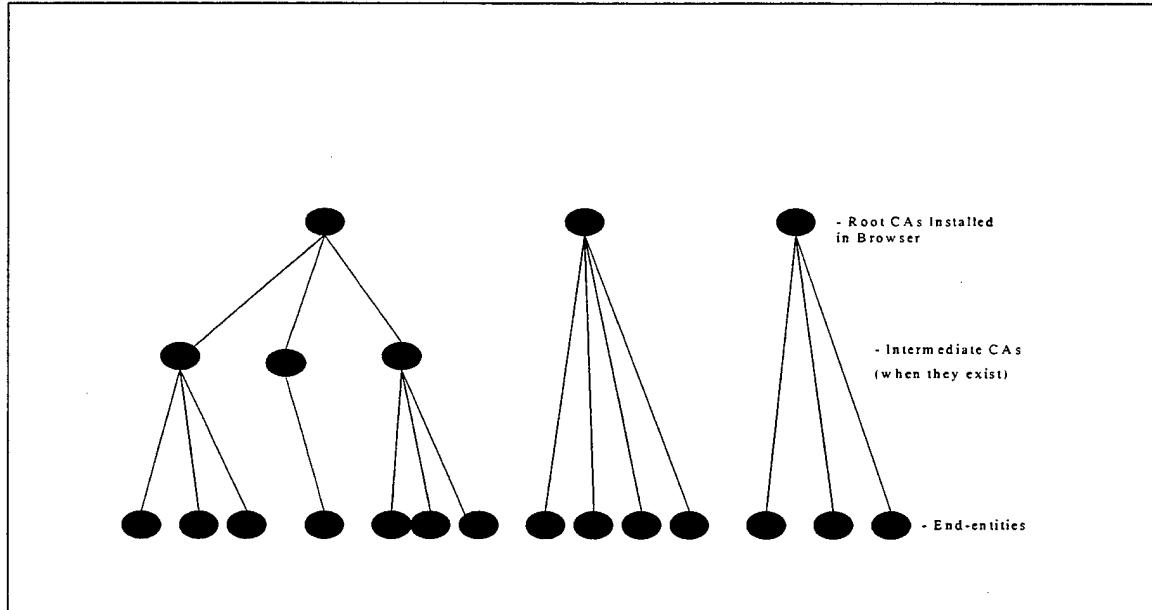


Figure 3, Web Model, (from Adams, 1999)

The Web model that Adams refers to has clear advantages in terms of convenience and simple operability. However, there are lots of security implications should be taken into when making deployment decisions for an environment. For

example, most browser users automatically trust the full set of pre-installed public keys, so security could be completely compromised if even one of those root CAs is "bad."

4. User-Centric Trust

The last model Adams refers to is user-centric trust. In this model each user is directly responsible and totally responsible for deciding which certificates to rely on and which to reject. Decision can be influenced by factors, although the initial set of trusted key often includes those of friends, family, or colleagues a given user knows. Adams states this model can be illustrated in the security software program Pretty Good Policy (PGP).

In PGP, a user builds (or effectively joins) a so called "web of trust" by acting as a CA (signing the public keys of other entities) and having his/her own public keys certified by others. When Alice receives a certificate reported online from Bob, she will see the certificate signed by David, who she does not know, but that David's certificate signed by Catherine, whom she does know and trust. For example, Catherine may have a certificate signed by Alice herself. Alice may then decide to trust Bob's key, by trusting the chain of keys from Catherine to David to Bob, where she may decide to reject Bob's key, judging that the "unknown" Bob is too many links away from the "known" Catherine.

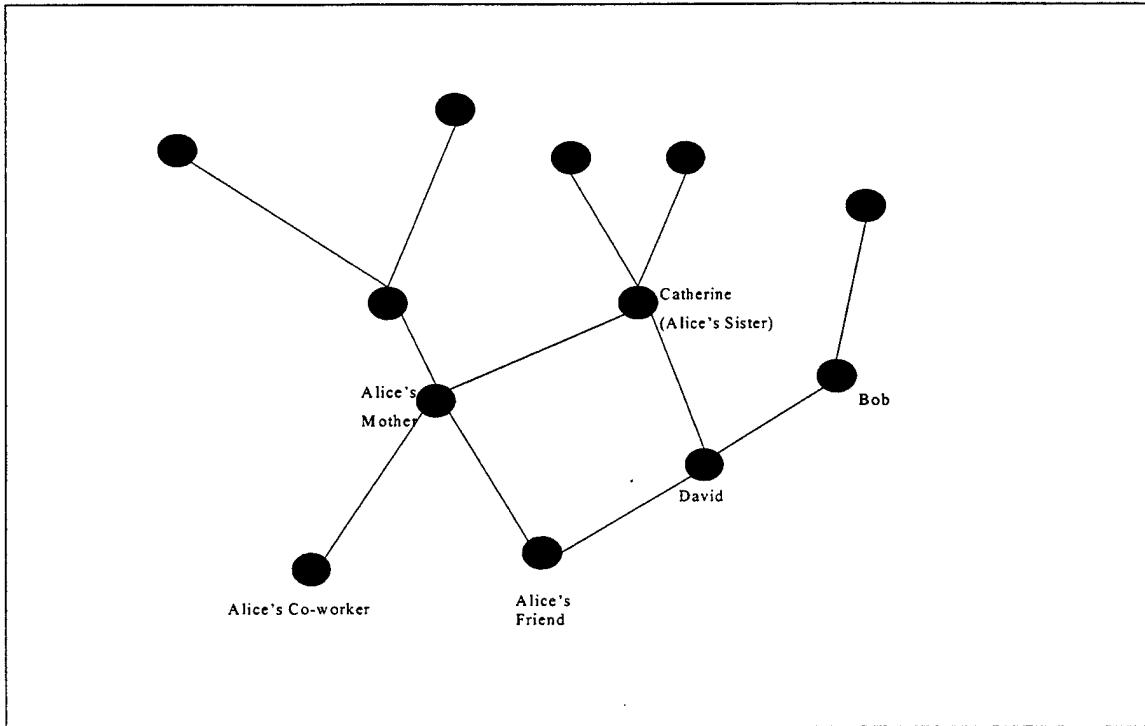


Figure 4, User-Centric Model, (from Adams, 1999)

Because of its reliance on user actions and decisions, the user-centric model may work in a highly technical and changing community, but it is unrealistic for the general, non-technical community, one which many users have little or no knowledge of security or PKI concepts. This model is most likely inappropriate in corporate, financial, or governmental environments because in these environments there is a desire to exercise some control over user trust; for example: such environments may want to enable or disable trust in a particular key or set of keys on an organizational basis.

C. RADIA PERLMAN

Doctor Radia Perlman wrote a paper entitled, *An Overview of PKI Trust Models*. The paper includes multiple models, some of which are similar to the ones mentioned

above. The models she describes are: A Single-CA Model, A Single-CA Plus Ras Model, An Oligarchy of CAs, Configured Plus Delegated CAs, Anarchy, Top Down, Up-Cross Down, Flexible Bottom Up, and Relative Names.

The Single-CA model consists of a single CA for the entire world. In this model, every piece of software and hardware would have to be configured with that CA's public key embedded in the firmware. The problem with this is that there is literally no organization that is trusted by all countries, companies, universities, military organizations, etc. She mentions that it is expensive and an inconvenience for organizations to obtain certificates from such a great distance and from an unrelated organization. When dealing with computer security issues, it is also good practice to change keys on a regular basis. In this Single-CA model, figure 5, if the CA were to change a key, the firmware in the entire world would have to be reconfigured and the CA's public key would have to be reconfigured.

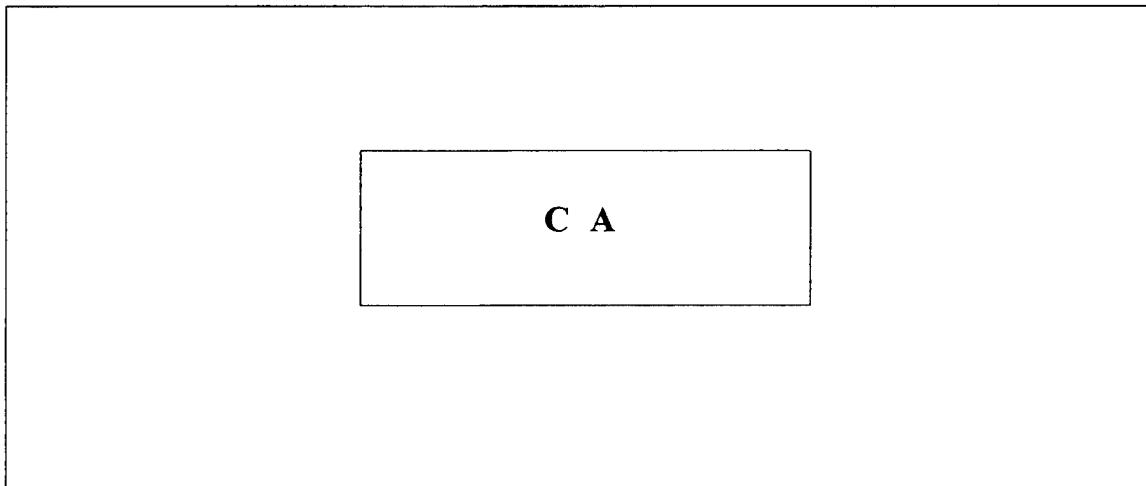


Figure 5, A Single – CA model (from Pearlman, 1999)

A Single-CA Plus RAs Model, figure 6, consists of a Single-CA and all the principals are configured with the CA's public key. The certificates are signed by the CA. Multiple registration authorities (RAs) are trusted by the CA to verify names and keys, and send a signed request to the CA. If the CA receives a valid signed request from the RA, it is granted a certificate. To most users, this model looks like the Single-CA model. One advantage of having multiple RAs is accessibility.

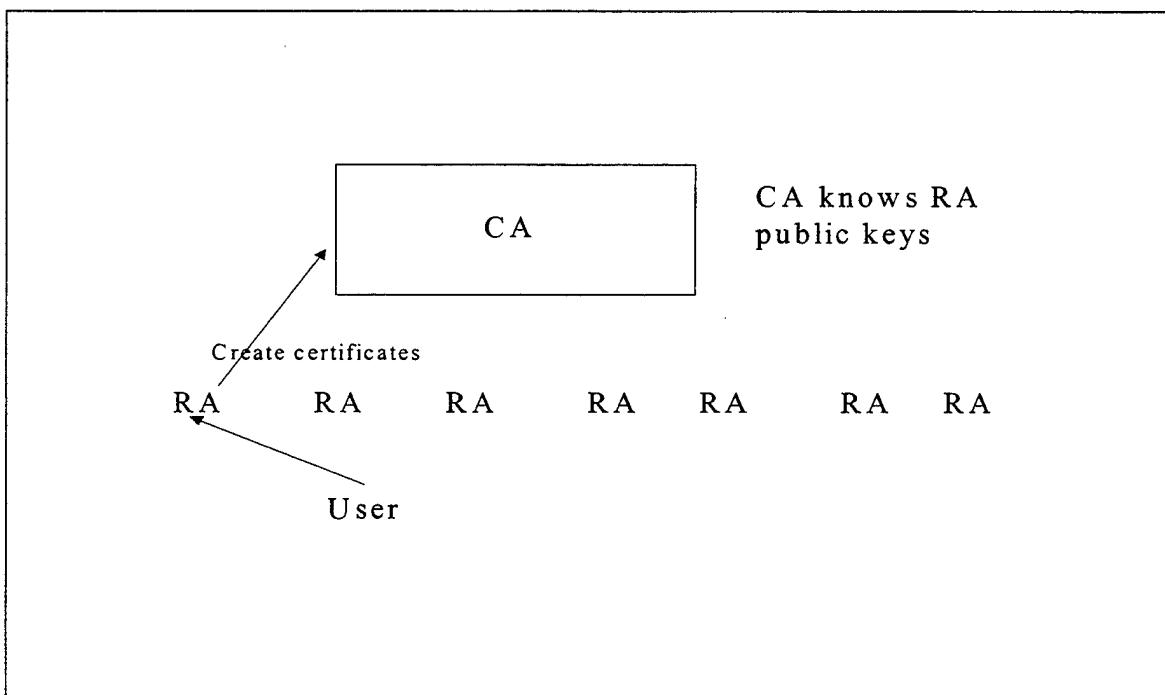


Figure 6, A Single – CA plus RA (from Pearlman, 1999)

In the Oligarchy of CAs Model, figure 7, everything is configured with a set of keys so there are perhaps dozens of organizations from which one can obtain certificates. This is similar to the Single-CA model where everything is configured with a public key of that Single-CA. The advantage of this model over the Single-CA is that competition

among trusted CAs should prevent abusive pricing for obtaining certificates. The chief disadvantage to this model is that it is less secure than the Single-CA model. Obviously, in the Single-CA model, if one key was compromised, all would be compromised. The security of the oligarchy model depends on all of the configured keys remaining secure. Compromise of any of the dozens of keys is as serious as compromise of the single key in the Single-CA model. Compromise is more than likely to be accomplished through a naïve user who might think that a workstation in an airport lounge or hotel room might be secure. It is possible to configure the workstation with bogus CA keys and impersonate the rest of the world to the naïve user of the machine. Changing just a small piece of information on the user's machine can make the machine attackable over the network. In theory, this is no different than installing malicious code on a publicly accessible machine. This is one problem that one can do little to protect oneself against.

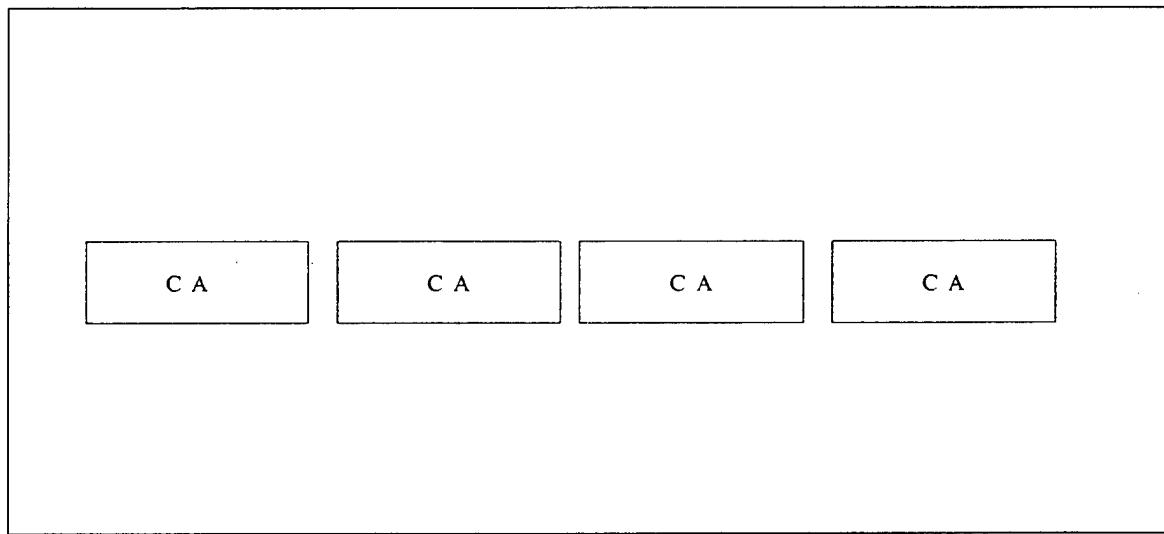


Figure 7, An oligarchy of CAs (from Pearlman, 1999)

The Configured Plus Delegated CAs model, figure 8, is implemented in current browsers, and is similar to previously mentioned models, the only difference being that the CAs whose keys have been configured into the users' workstations can sign certificates authorizing other CAs to grant certificates. Configured CAs, keys configured by the CA, and delegated CAs, keys that have been delegated and authorized by those CAs to act as a CA, are completely trusted and any certificate from any of those CAs will work. An advantage to this model is that it allows users to obtain certificates from more places, since they no longer need to get a certificate directly from one of the configured CAs. Instead, they can obtain a certificate from a delegated CA.

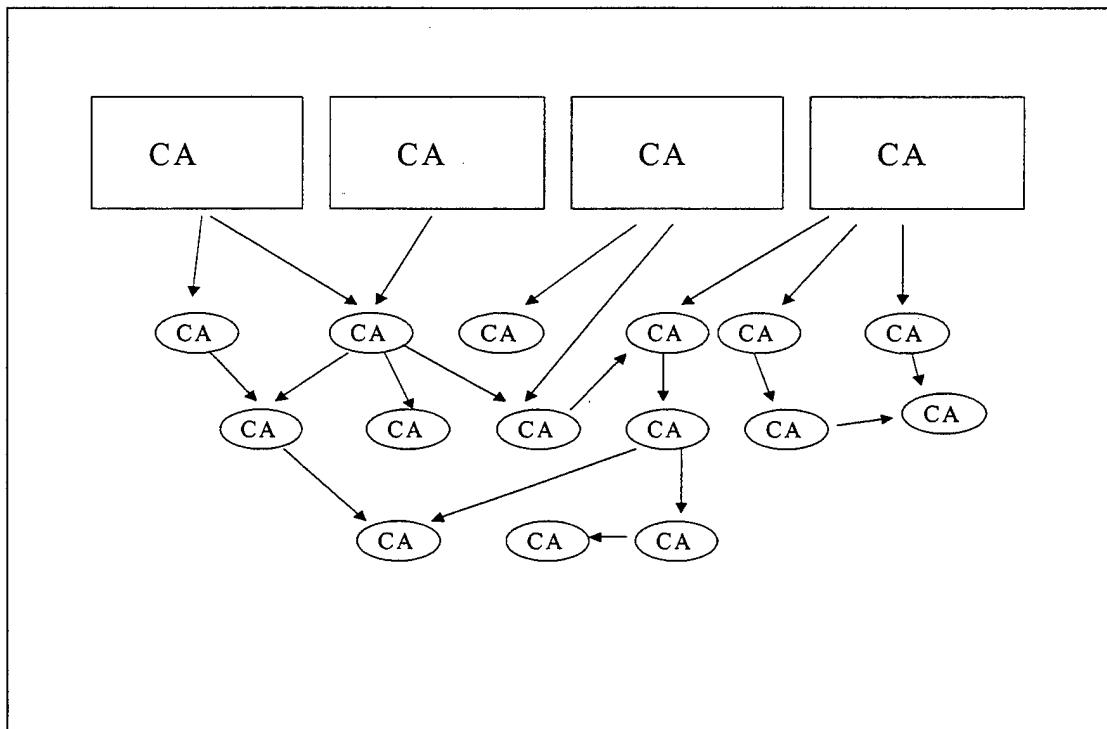


Figure 8, Configured plus delegated CAs (from Pearlman, 1999)

This model is somewhat more secure and convenient for users to obtain certificates. This is largely because there is a significant increase in the number of CAs, it is likely that users like Alice can find a CA sufficiently nearby that she can physically visit, making it more convenient to give the CA Alice's public key. However, if a delegated CA is compromised in theory, the certificate authorizing that CA can be revoked. But the compromise may go unnoticed. And unfortunately, it is not uncommon for the PKIs to ignore the revocation issue. In short, this model can even be less secure because there are more CAs, and theft of any of their keys can enable the thief to impersonate anyone to anyone. In addition, with delegated CAs, the certificate chain is no longer, and verification is less efficient.

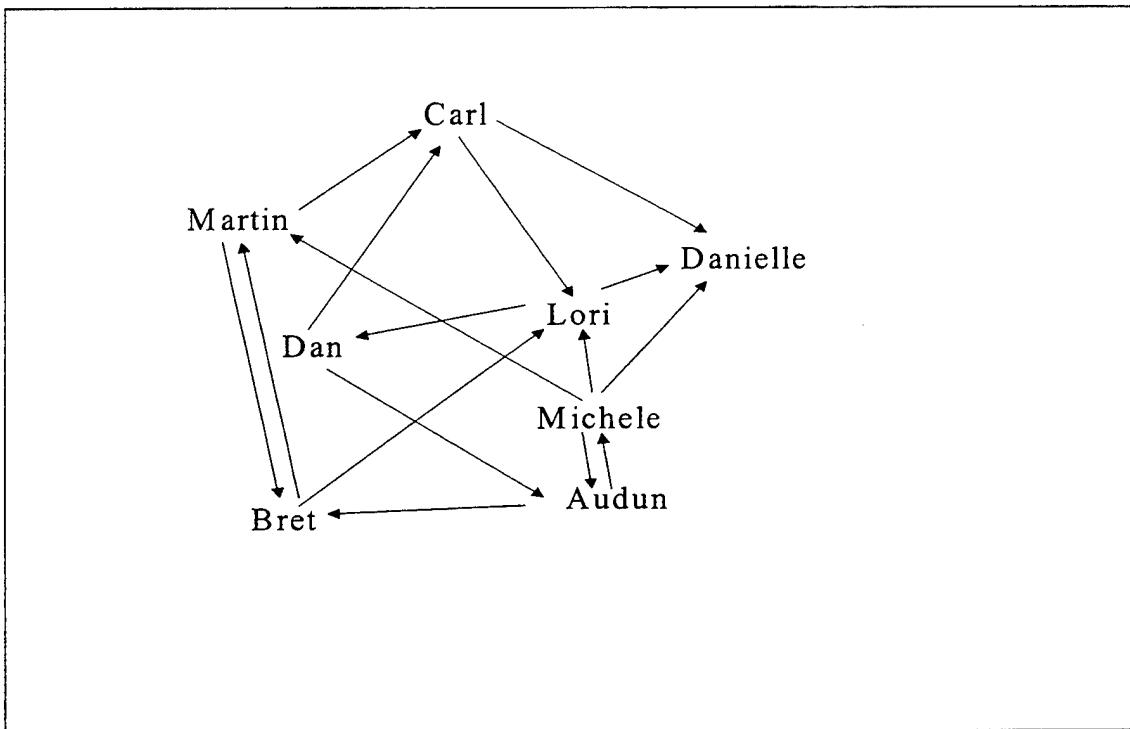


Figure 9, Anarchy (from Pearlman, 1999)

The Anarchy model, figure 9, is used by the original public domain PGP. In this model, each user starts off by configuring public keys they have securely learned out of band. They then obtain certificates through a number of other means, such as e-mail or downloading certificates from public databases. In PGP, if you know the person that sends an e-mail, you sign a certificate for them. This does not scale beyond a relatively small community of trusted individuals. If PKI was the security choice of the internet, how big would the database have to be for a conservative estimate of approximately 100 million Internet users, each with about 10 certificates: a database of more than a billion certificates which would make searching for a certificate difficult. Even if you did find a path mathematically created in a chain from a key you trusted earlier, how would you know whether you could trust that chain to that key again? In PGP the certificates are only verified by the signer who identifies the subject at hand. Trust is considered a local matter. However, there is no way of judging the trustworthiness of someone several links down in the chain whom you have never met. The trust of information in the public domain PGP only applies to the first link in the chain. So, to close the anarchy model, there is no core set of configured CAs. Instead, each individual starts with a personally configured set of trusted public keys.

In the Top-Down Model, there is only one configured root CA key, and that CA can delegate to other CAs, which can delegate to other CAs. The model consists of a hierarchical namespace, and a CA is only trusted to certify name-to-key mappings for names in the subtree of the namespace with the root being that CA. This rule of trusting a CA only for a portion of the namespace is known as name subordination and makes this

model much more workable than the previous models where the CAs were completely trusted to certify names.

In the top-down model, each user starts out knowing the root key, and retrieves all the certificates from the root down to their own key. For example, Alice can authenticate to Bob by sending him all the certificates from the root down to herself. Since Bob also starts out knowing and trusting the same root key, he has a path to Alice's key.

In order to play the role of a CA in this hierarchy, it is necessary to follow a set of policies defined by someone other than the CA. In other words, security would be one size fits all and every organization would have to be equally careful. Top-down models have the disadvantage that the entire PKI depends on the security of the single root key.

To change a root key would require a massive reconfiguration of everyone.

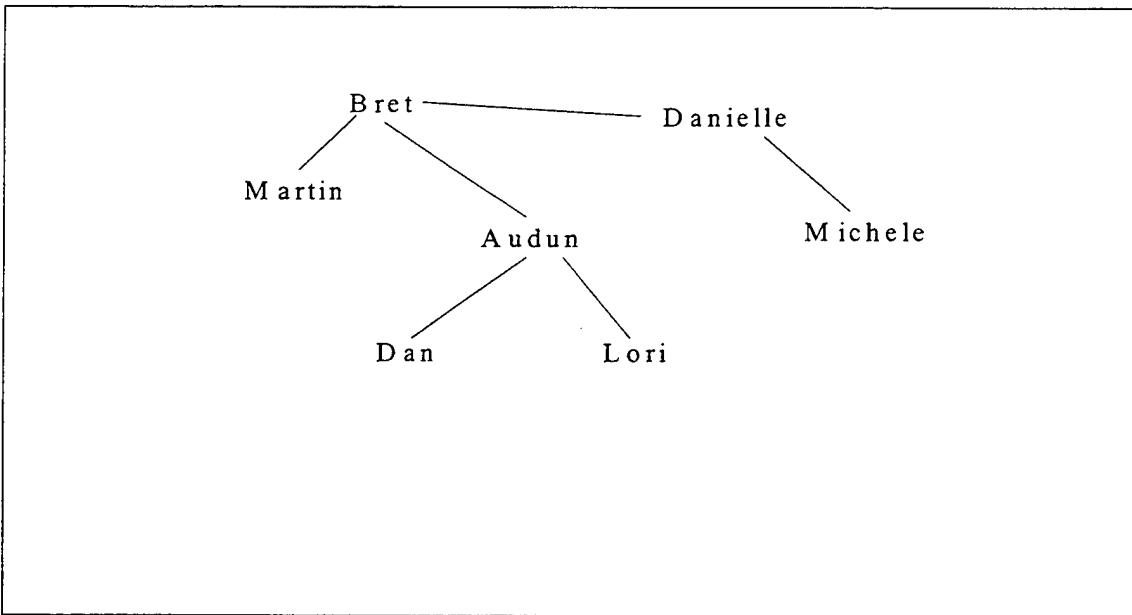


Figure 10, Up-cross-down (from Pearlman, 1999)

The Up-Cross-Down model, figure 10, assumes a hierarchical namespace with a directory structured so that any name can be found. It also assumes the name-subordinate rule. Each node in the namespace is represented by a CA and it contains three types of certificates.

- Down: a parent certifies the key of the child
- Up: a child certifies the key of the parent
- Cross: any node certifies the key of the other

The basic rule for this model is you go up as many times as necessary to reach a cross certificate, or at least a common ancestor, then you go across at most once, and then down to the target name.

This model has many advantages over the models described in the previous sections. Security can be deployed within an organization without the need to obtain certificates from any outside organization. No single compromised key will cause a massive destruction of the local system. There is no preordained root organization. The entire PKI can consist of independent intra-organizational trees loosely coupled with cross certificates. And, security within one's own organization is completely within one's own hands. If the organization manages the key well, compromised keys outside the organization will not affect what are presumably the most security-sensitive operations, namely authentication between users of the organization. This is primarily due to the fact that the path between users in the organization stays inside and does not go outside the organization.

In the Flexible-Bottom-Up model, use of a Public Key Infrastructure X.509 (PKIX) extension known as name constraints, certifies a CA, but only for the explicitly named portion of the namespace. The name constraints field can include permitted names, excluded names, or both. When evaluating paths, it is necessary to follow every link for which the name constraints are still included as the target name. This model will have approximately the same search complexity as the anarchy model. Every certificate link must be explored because it might lead to the target name.

The flexible bottom-up model has the disadvantage that it allows unscalable structures to be created, but it has the advantage of certain flexibility not accommodated in the up-cross-down model. One example where more flexibility might be desirable is to support a mesh of root CAs rather than a single root CA. To support this structure, the usual up-cross-down default-name constraints would be set in all certificates except those of the mesh of roots. This model has all the advantages of the up-cross-down model, provided it is not abused so much as to make searching for paths intractable, and it allows more flexible trust rules than strict up-cross-down.

IV. D-M MODEL

A. ISSUES OF TRUST

There are many "models" available for review on the World Wide Web (WWW), many of which have similar names and similar definitions, but yet the models differ. What we have discovered is that there is no standard for how we measure trust. In particular, there is no model which incorporates both mandatory and discretionary trust. In this chapter, we will introduce our model and describe the components involved in our model. The motivation for using our model in the context of trust stems primarily from the appropriate use of discretionary and mandatory trust policies in organizations to ensure precision, consistency, and added assurance in trust. While it is appealing to approach trust intuitively, it is important to recognize that many good ideas turn out to be not so good on close inspection. This chapter introduces the D-M Trust Model, which provides for representing both discretionary and mandatory policy about trust. It is important to keep in mind that mandatory policy dominates discretionary policy. The D-M model is built on top of, or is intended to be used in conjunction with existing computational models of trust.

Audun Jøsang, who has written extensively on the subject of trust and trust models, developed a trust model that utilizes an opinion model and a process he calls subjective logic that consists of a set of algebraic operators. His model can be applied to a number of applications including authentication, security evaluation, artificial intelligence, and e-commerce. However, Jøsang's model and that of other researchers do not bar the specification of both discretionary and mandatory policy about trust. The real

value of the D-M model, is that it addresses the need to model both of these types of policies explicitly and concurrently.

B. TRUST SYSTEMS

Trust systems are no more secure than trust in people. An example of this is trusted users. There are many concerns with users and establishing trust models suited for both corporate and defense organizations. To start, one must be cautious about user registration. When a user, or entity, becomes an authorized member of a security domain, a certain level of trust will need to be established. This trust must be established for an audit trail and stored in a database, well before the creation and or exchange of initial material such as shared passwords or personal identification numbers (PIN) by a secure, one-time technique.

Regardless of the composition of a trust system or trust model, issues can arise in the discussion of such models. For example, do we try to model trust at the user level where the user is an individual, or should we model trust at the user level, where the user is composed of a group of two or more people? Absence of complete trust between two entities is an issue in a number of areas:

- intra-organizational trust – trust between sub-organizations within the same organization
- inter-organizational trust – trust between organizations
- trust in organizations – trust between organizations and those they serve
- social trust – trust between individuals in particular communities

C. D-M MODEL

The purpose of the D-M model is to represent policy while at the same time understanding the need for distributed decision-making. A policy is a statement of a definite course of action intended to influence the actions of constituents; a definition of roles, obligations, and the right of constituents. A procedure is an implementation of policy as rules and ways to make exceptions to rules. Procedures can be either manual or automated. The real value in the model is that it combines the models mentioned in chapter 3, and allows top-down, bottom-up, and lateral flow of both information and trust while allowing decisions to be made at the lowest possible levels in a hierarchy of actors.

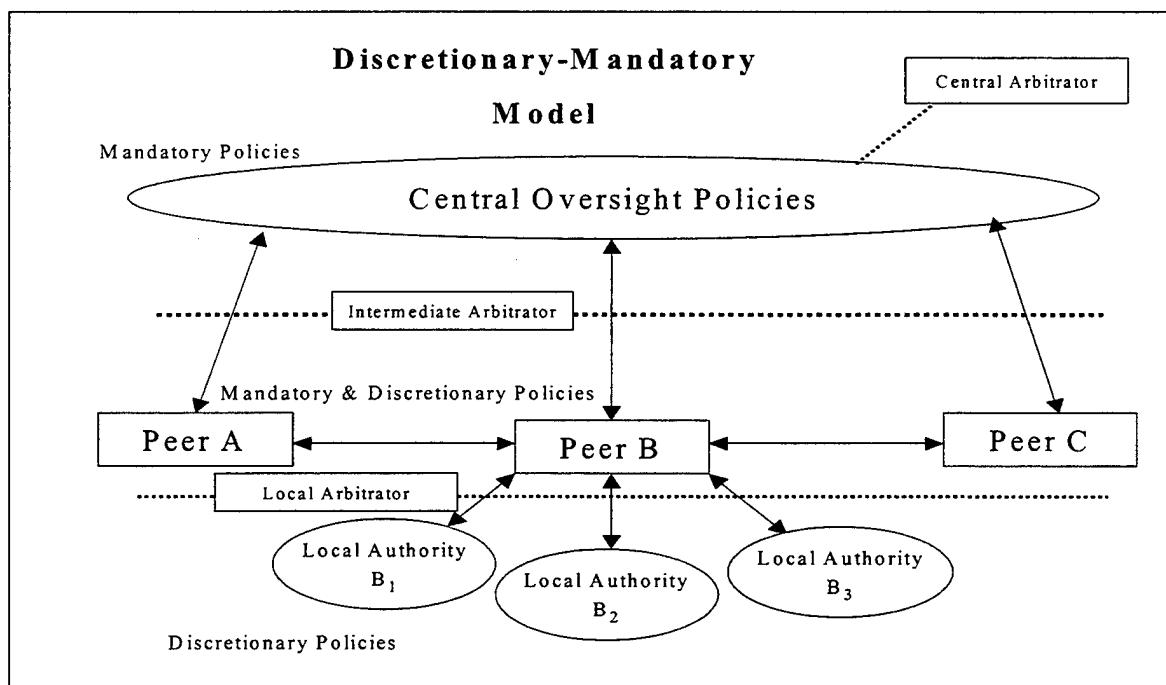


Figure 11, D-M Model

D. THE BASIC COMPONENTS OF THE D-M MODEL

In order to represent the complexities of the D-M model, we posit the need for basic terminology of:

- Actors
- Groups
- Levels of Arbitration
- Policies

1. Actors

Actors are the primary manipulators that can change the state or structure of the system. In the D-M model, peers are also actors at the same level in the hierarchy of an organization. Actors are not necessarily humans; an actor can be human, computer process, or some combination of the two. For example, an actor can be comprised of any part or a military organization, government, or civilian company. Actors can be broken down further into individual tactical units such as a regional CINC, or Joint Task Force (JTF) for the Navy, or battalions and fire support units, for the Army.

There are many times when discretionary trust is required between peers or actors. For this reason our model will incorporate a flexible-bottom-up scheme also referred to as cross connect, allowing end users discretion and flexibility in making their own decisions.

2. Groups

Groups are represented by the large oval circle we term ‘central oversight policy.’ These groups are not specific to any one entity. They could also include but are not

limited to: chains of peers, chains of local authorities, and groups of central policy-makers. This is key to allowing for a distribution of models to be chained together. Contrary to a strict hierarchy model, in which all of the actors in a designated community trust a single root agency, our model includes a distributed trust architecture where trust is distributed among two or more groups controlled by an authority (e.g., a government) to oversee the subordinate groups. The D-M model is based on a hybrid distributed trust architecture, which may be referred to as a fully peered architecture, because all the actors, or peers, are independent actors.

There is no central node or root group in the D-M model. Groups are peers with each other. An example of the use of this construct is sets of multiple interconnections of independent agencies from different organizational domains.

3. Levels Of Arbitration

In order to provide control, an arbitrator must be established to regulate, monitor, and oversee all policies associated with trust in that part of the model. The central arbitrator has the final say if a dispute occurs between organizational chains, dictated by the Central Oversight Policies. For example, if the regional CINCs have a dispute over mandatory policy, or in this case – doctrine, then a Central Arbitrator, that is, Chief of Naval Operations (CNO), would be called on to act as a mediator between the CINCs. We have also included an Intermediate Arbitrator to oversee the actors; this is necessary for resolving questions or disputes that arise over the use of both mandatory and discretionary policies. The Local Arbitrator is a safeguard between the actors and the various local authorities that could be imposed in a given system. For example, a

question might arise as to which policy should be followed when two or more related policies are in conflict (i.e., inconsistency).

4. Policies

All formal organizations have policies. As we stated earlier, a policy is a statement of a definite course of action intended to influence the actions of constituents; a definition of roles, obligations, and the right of constituents. Webster's dictionary defined it as "a definite course or method of action selected from among alternative and in light of given conditions to guide and determine present and future decisions."

(Webster, 1987) Explicit representation of policy is desirable so that actors can maintain, enforce, and reason about policy. A formal representation of policy is desirable because such a representation is amenable to some level of automation: computer-based tools can be used to assist the users of policy in maintaining, enforcing, and reasoning about policy. (Sibley, Michaiel, and Wexelblat, 1992) Our primary strategy for assuring trust is to provide a model that simultaneously incorporates both discretionary and mandatory policies.

a. Discretionary Policy

Discretionary policies are so named because they permit the actors in any given organization to decide which policies are to be enforced at their own discretion. An example of discretionary policy is one that is seen in every United States Naval War Ship. It is a position held by a Tactical Action Officer (TAO). The TAO is considered a watch station that is manned 24 hours a day in Combat Information Center (CIC). The TAO is given weapons release authority for safe and accurate deployment of the ship's

offensive and defensive weapons from the Commanding Officer. The TAO's role is based on Discretionary policy in that the TAO can fire on the enemy or assume a defensive posture without the Commanding Officer's permission.

b. Mandatory Policy

Mandatory is the inverse of discretionary. Mandatory policies must be enforced by the system – this type of policy is not at the discretion of the actor. A Mandatory policy dominates discretionary policies. To further explain what we mean by dominates, consider the example of the TAO. The TAO is given the inherent responsibility to use Discretionary policy as it relates to offensive and defensive operations of the ship. The TAO ‘may not’ deviate from Mandatory policies such as specific Rules of Engagement (ROE) that are to be strictly enforced. The D-M model references a central oversight policy that is mandatory in nature. For example, specific theatre ROEs are Mandatory policies that dominate any Discretionary policy under the TAO’s control, designed to provide further guidance and to not create a National incident.

In practical matters, the choice between making a policy mandatory or discretionary can be tied to the risk of incurring a loss and the magnitude of the loss due to the violation of policy or lack of uniform enforcement of a policy.

c. Interaction Between Policies

In systems that incorporate both mandatory and discretionary policies, the discretionary policy serves to provide fine granularity within (but cannot substitute for) the mandatory policy. For example, our military need-to-know security policy in which

each actor has a responsibility to determine that another actor has a valid requirement for information, even though the other actor has a clearance for the information, is a common discretionary policy.

E. LIMITATIONS OF D-M MODEL

The D-M model has some known limitations. In the fast pace of the electronic cyber world, it is at most, difficult to derive a model that covers every situation. There are always going to be challenges, especially since the design architecture of our D-M model is based on a distributed system. This distributed architecture may be referred to as a fully-peered architecture, because all the peers are independent peers. The use of separate processors for databases in the peer-to-peer relationship is hardware intensive. Even if no processors were involved totally separate organizations are involved, the capacity of those organizations to have the same set of mandatory rules as the other is going to be difficult to enforce, albeit followed. Secondly, we have stated that the actors are trusted to use their discretion and they are trusted to make and pass on decisions. Improper implementation of the model could possibly allow the actors to process high volumes of data that could potentially take up unnecessary bandwidth, or unnecessary channels that could potentially slow the entire system down.

Another limitation of this model is in that it allows multiple structures to be created. By allowing users to create their own networks of trust, the size of distribution of trust will grow in incremental proportions, possibly to the point of being out of control with so many trusted networks. On the other hand, the model provides some degree of flexibility by permitting the end user to make decisions on which they will trust. This

model has all of the advantages; provided it is not abused so much as to make searching for paths intractable, and it allows more flexible trust rules than strict up-cross-down model.

F. SUMMARY

The centerpiece of our work is the importance in ability to combine discretionary and mandatory policies. These policies specify how to manage, protect, and distribute information. As stated earlier in the chapter, the real value in the D-M model is that it allows top-down, bottom-up, and lateral flow of both information and trust while allowing decisions to be made at the lowest possible levels in a hierarchy of actors. In the following chapter, we shall develop a case describing the use of the D-M model and apply the model to a battle group connectivity scenario of an American aircraft carrier steaming in the Persian Gulf.

THIS PAGE INTENTIONALLY LEFT BLANK

V. CASE STUDY

A. CASE STUDY - AEGIS PLATFORM

The Aegis platform uses the Joint Integrated C4 System to maintain command, control, and coordination links with the Joint Task Force (JTF) commander and the various combatants and support forces JTF has assigned to the area of responsibility. In this example, the most senior Commanding Officer afloat would fulfill the role of local authority and at the same time could possibly act as a local arbitrator in situations in which there exists discussion or confusion about discretionary policies between the actors. The JTF commander could act as the intermediate arbitrator (see figure 11). It plays a very important role in acting as a mediator of discretionary and mandatory policies between the actor/peer (Aegis platform) and the Area of Responsibility(AOR) Commander in Chief (CINC). If for some reason there was some conflict of orders regarding the mandatory policy, the AOR CINC could fulfill the role of Central Arbitrator.

Additionally, the Tactical Data Exchange System and the Tactical Command Information System promulgates mission data updates and processes command and control information concerning the actual missile-fire missions. Without getting into the classified details of the entire missile launch process, this is an example of mandatory policy specifically instructing the vessel in the requirements of a missile launch fire mission. The Aegis platform takes the assigned mandatory policy used to launch the selected weapon system, and gives the commander of the vessel the inherent ability to use discretionary policy to get to the assigned missile launch area.

The chain-of-command influences and standardization of processes and outputs become more critical as the Aegis platform maneuvers into a missile-launch position to execute the assigned mission. For this reason, overall there is less need for horizontal coordination of trust between units / actors because the coordination, is already in place and a lesser need for localized direct supervision at the lower levels of the organization.

As the mission process reaches the culminating point at the lowest level, the interdependencies between the Aegis Missile Platform and its Combat Systems Department become pooled in nature, thus acting as a cooperative trustworthy machine. This is a result of the D-M model which has incorporated a flexible bottom-up scheme also referred to as cross connect, allowing end users discretion and flexibility in making their own decisions.

At this stage, the focus is on the actual weapons release of the mission and execution follows an exacting set of procedures that does not allow for any deviation or margin of error (mandatory policy). Thus, trust will have been established throughout the entire Aegis ship (discretionary policy), which also does not encourage error, thus allowing end users or actors, to make critical decisions about weapons release. The impact of chain-of-command influences is most likely at its highest level at this stage in the control and coordination of the mission, with the Tactical Actions Officer (TAO), who has assumed the responsibility of role-based trust to fight the ship, is working hand-in-hand with the Commanding Officer. Underlying and facilitating the interactions of all of these coordination techniques is a common doctrine shared by all echelons in the organization normally called “Commanding Officer’s Battle Orders.”

The TAO's role, based on trust, represents an important form of presumptive trust found within organizations. Role-based trust constitutes a form of personalized trust based on the predicated knowledge that a person occupies in a particular role in the organization rather than specific knowledge about the actor's capabilities, dispositions, motives, and intentions. The Navy achieves this by sending officers to Department Head school. This is where they learn the intricacies of weapon systems and how to apply those weapon systems to "fight the ship," as we say in the Navy.

Thus, to the extent that actors within an organization have confidence in the fact that role occupancy signals both an intent to fulfill such obligations and competence in carrying them out, as in the case the TAO, individual actors can adopt a sort of presumptive trust based upon knowledge of role relations, even in the absence of personalized knowledge or history of prior interaction.

Such trust develops from and is sustained by the actor's common knowledge regarding the barriers to entry into organizational roles, their presumptions of the training and socialization processes that a TAO, or role occupants undergo, and their perceptions of various accountability mechanisms intended to ensure role compliance and trust in fellow actors in the Chain of Command.

As in the previous instances, the Aegis platform is representative of Actors, Groups, Levels of Arbitration, and Policies of an entire organization of trust and can be applied to the D-M model. Thus the D-M model can be used to reference any structure or organization to include adequate supporting staffs to aid personnel supporting the organizational concept of operation, by utilizing a combination of mandatory and

discretionary policies. As the lead actor of the entire organization, the skill sets, knowledge sets and work capacity of its actors are presumed to be of a high level of trust. The following case study will show not just one platform, as in the Aegis example, but an entire Joint Task Force Battle Group that includes two foreign naval vessels. The logic in proposing a battle group connectivity case study is to show how the D-M model can be applied for a relatively large and dynamic organization. We use the term ‘dynamic’ here because ships and aircraft can join and leave the battle group at any time.

B. CASE STUDY – BATTLE GROUP CONNECTIVITY

Let us look at a case of an American aircraft carrier steaming in the Persian Gulf conducting normal flight operations. It has in company an American Aegis cruiser and Fast Frigate, along with a British destroyer and a Dutch frigate.

The Dutch frigate acquires radar contact on an unknown aircraft traveling inbound which it classifies as hostile and transmits the track to the rest of the Battle Group. The Dutch frigate then loses radar contact with the aircraft but continues to track it as hostile in the Battle Group database.

At the same time as the loss of radar contact by the Dutch frigate, the unknown, potentially hostile aircraft, is acquired by the Aegis cruiser at a distance of 100 kilometers from the aircraft carrier. The Aegis system determines it is the same unidentified contact classified as hostile by the Dutch frigate. It is within the air launched weapons envelope of multiple theater threat aircraft. What should the Aegis cruiser do?

Although U.S. doctrine and the standing rules of engagement would likely allow the Aegis cruiser to destroy the unknown aircraft, that would make little difference in

world opinion if the aircraft turned out to be a passenger jet. Alternatively, if the Aegis does nothing, and the aircraft turns out to be an attack aircraft that launches its weapons on the carrier, the cruiser will have failed to carry out its duty as a naval tactical warship.

The commander must decide how much trust to place on the information coming from an Allied Dutch frigate. If there is an established relationship over time, common procedures and training to establish trust among the two platforms, then the Aegis cruiser can act with confidence on the data provided by the frigate. However, if there are no commonalities and no history of an established trust relationship, then the trust factor for this individual piece of data will be low.

An interesting side note arises when working with allies. One should question and take into account the Operational Security (OPSEC) of the given situation. Consider that the United States is at war and has many Allies in the Coalition. In such a case not all coalitions are created equal. For example, should we place the same level of trust in, say the Syrian allies, as the U.S. would place in the British allies? The answer to this question may be obvious, but diplomatically it is not an easy answer. One has to be careful of the type and amount of information that is disseminated to, and received from Allies along with one's trust in the information. Not all Allies are created equal.

Although this is an example of a Joint Task Force with several allies, the problem of trust does not solely rest on the shoulders of the foreign navies. If one examines the American Aegis cruiser, one will find one of the most technologically advanced warships in the world. In addition, in company with the cruiser is another actor, an American Frigate, which has a combat systems suite that is dwarfed in the shadow of the Aegis

platform. Nonetheless, information is still disseminated into the Combat Direction System, for all to access. The problem, which one might encounter, is that of the age of technology between the two American platforms is significant and the warriors on the new, highly technological sophisticated ship may not trust the data disseminated from outside actors.

If one were to look back to July, 1988, one could use the case of the USS Vincennes as an example. The USS Vincennes shot down a jetliner because the Commanding Officer trusted an information system that mistakenly identified the jetliner as a hostile target. With the TAO and Commanding officer receiving not only information across the tactical data links, but information from their own operations specialist second class petty officer in the Combat Information Center. The specialist had earned at a great deal of responsibility and had been given a great deal of discretionary trust, which allowed him to make decisions, for the entire chain of command for the position he was in. The Specialist role is to evaluate the accuracy of the tactical data displayed before him. The Specialist then passes the tactical information up the chain of command for further review and action if necessary.

Even though no formal rules of engagement were broken, this incident resulted in a political crisis for the United States as well as adding to tensions in the Middle East. In the vicinity this incident was the American Frigate, USS Sides, who identified the unknown target as a jetliner. When the USS Sides announced that the perceived unknown hostile air contact was a jetliner, their information was promptly dismissed. This in part because of a combination of human factors of trust the USS Vincennes'

Commanding Officer placed on his weapons system for its accuracy and precision, the newest to the fleet at this time, to identify and evaluate the unknown air threat as a hostile platform. With increasingly automated systems coming of age, local decision makers must understand how to assess the context of a given situation and evaluate trust in these systems and be given the discretionary permissions to make their own judgments rather than following a predetermined strict decision sequence.

Failure to evaluate all disseminated data and the amount of trust in a given piece of information resulted in tragedy. However, if a system were in place that incorporated the D-M model that allowed for feedback from individual actors, for example the USS Sides, as well as discretionary decision-making ability to be negated only by mandatory policies from higher echelons in the command structure, i.e., the JTF or Intermediate Arbitrator, this tragedy could possibly have been avoided.

C. TOOLS TO ASSIST NAVY SHIPBOARD DECISION MAKERS

1. The Rapid Anti-Ship Cruise Missile Integrated Defense System

The Rapid Anti-Ship Cruise Missile Integrated Defense System (RAIDS) is a tactical decision aid for the Commanding Officer/Tactical Action Officer and the Electronic Warfare Supervisor. It provides automatic display of anti-ship cruise missile (ASCM) threats, depicts active and passive sensor displays, and shows the status of existing terminal self-defense systems. RAIDS, a multiple microprocessor-based system, considers threat capabilities, environmental data, electromagnetic interference data, own-ship maneuvering parameters, and approved tactical doctrine to develop a dynamic tactical decision matrix that provides a ship's anti-ship missile defense (ASMD)

coordinators with concise and real-time tactical engagement recommendations. These recommendations are continually and automatically evaluated for effectiveness and updated as appropriate. It is displayed to the TAO as a decrementing timeline, providing a recommended firing solution for the potential hostile threat (RAIDS, 1998).

2. Ship Self-Defense System (SSDS)

Ship Self-Defense System (SSDS) is a combat system that intends to integrate and coordinate all of the existing sensors and weapons systems aboard a ship. The system will eventually be installed aboard most classes of non-Aegis ships. SSDS makes it possible to automate the detect through engage sequence using identification and engagement doctrine statements. SSDS is one of the first actual implementations of Network Centric Warfare (NCW) and if implemented with the D-M model, will add increased value of trust throughout the Naval Fleets. (SSDS, 2000)

a. *The Principal Threat*

The principal air threat to U.S. naval surface ships is a variety of highly capable anti-ship cruise missiles (ASCMs). These include subsonic (Mach 0.9) and supersonic (Mach 2+), low altitude ASCMs. Detection, tracking, assessment, and engagement decisions must be accomplished to defend against these threats, with the duration from initial detection of an ASCM to its engagement with weapons, typically on the order of a minute or less with a certain trust level. SSDS is designed to place a great deal of discretionary trust in sensor input and user recognition to accomplish these defensive actions.

b. System Description

With radars and anti-air weapons for self defense of today's ships and aircraft carriers installed as stand-alone systems, considerable manual intervention is required to complete the detect to engage sequence against incoming missiles, or ASCMs. The Ship Self Defense System (SSDS) is designed to expedite that process. SSDS, consisting of software and commercial off-the-shelf (COTS) hardware, integrates radar systems with anti-air weapons, both hard kill (missile systems and rapid fire gun systems), and soft kill (decoys).

SSDS integrates previously "stand-alone" sensor and engagement systems for aircraft carriers and amphibious warfare ships, thereby supporting the Joint Vision 2010 concept of full-dimensional protection, by providing a final layer of self protection against air threats for individual ships. By ensuring such protection, SSDS contributes indirectly to the operational concept of precision engagement, in that strike operations against targets are executed from several of the platforms receiving SSDS (SSDS, 2000).

3. Cooperative Engagement Capability

Cooperative engagement, also referred to as sensor netting, will allow large numbers of CEC-equipped surface ships and aircraft to operate as a single "distributed" air-defense system capable of passing fire-control-quality radar target measurements in real time across the entire force. It is another example of what has been coined "Network Centric Warfare." The CEC system features two primary components--a cooperative engagement processor (CEP) and a data-distribution system (DDS), which acts as the

CEC communications relay--and a series of modifications to already-fielded combat systems.

In CEC operations, radar measurement information on airborne targets from shipboard air-search radars is provided to the CEP, which reformats the data and sends it to the DDS. The DDS then encrypts and transmits the data to other ships participating in the CEC network (referred to as CUs – Cooperative Unit). In a fraction of a second, the DDS receives all other CU data and forwards it to the CEP. The CEP combines all of the unprocessed sensor-measurement data into an identical air picture--consisting of continuous composite tracks of all targets. The same picture then is available for display and use by each individual platform's sensor and engagement systems. The DDS uses a narrow directional signal that is highly resistant to jamming and/or hostile intercept, and that allows simultaneous unit-to-unit communications between and among the various participating units (PU), permitting the DDS output to be used as real-time fire control data. These data are passed to the ship's combat system as fire-control-quality data that the ship can use to engage targets without actually tracking them with its own radars.

The CEC takes full advantage of the diverse range of capabilities achievable by the participation of multiple ships equipped with multiple types of sensors throughout the operating area. Combining the varying sensor inputs available synergistically enhances the completeness of the common CEC data picture--and thereby enhances the ability of the CEC-equipped ship to track and destroy incoming ASMs. CEC provides a capability, referred to as "engage on remote," whereby a ship that does not originate the tracking

data can launch missiles at targets within the weapons range identified in the CEC composite track picture.

D. SUMMARY OF KEY POINTS

The purpose of this chapter was to show the importance in the ability to combine discretionary and mandatory policies from a sensor to shooter perspective on an Aegis platform and that the D-M model could be used for most any situation in most any organization. We realize there is a need for a computational model and that models have to provide various functions tailored to the specific organization. As stated earlier, the real value in the D-M model is that it allows the core principles of top-down, bottom-up, and lateral flow of both information and trust while allowing decisions to be made at the lowest possible levels in a hierarchy of actors. This was shown by the lack of trust between the two United States Ships. If in fact the USS Vincennes would have assessed the information from the USS Sides, utilizing the core principles of the D-M model, the passenger liner may not have been shot down, and a world crisis could possibly have been avoided. In addition, if the JTF could have been acting as the Intermediate Arbitrator, as in the D-M model, they could act in the capacity of the arbitrator and settled any discrepancy in policy and data flow.

We would also like to surmise that the implementation of Network Centric Warfare is starting to develop in the areas of SSDS and CEC. Network Centric Warfare (NCW) has emerged as a new concept for the US Navy. NCW capitalizes on technology to obtain and maintain an enhanced situational awareness and uses a distributed firepower of the collective force to fight in a battle. This distributed firepower is the key to the core

principles of top-down, bottom-up, and lateral flow of both information and trust while allowing decisions to be made at the lowest possible levels in a hierarchy of actors in the D-M model. Thus the D-M model is not just the beginning of NCW, it is better!

VI. CONCLUSION

A. SUMMARY

Trust has been defined by researchers in many different ways. In fact, if one examines the many definitions, one might come to the conclusion that existing trust models are an amalgamation of different beliefs and ideas. The authors feel that trust is a cognitive function and modeling trust is an attempt to emulate the way a human assesses trust. There are a number of trust models that represent attempts to define and assign metrics to trust. These models address the notion of trust in many different ways and their definitions and metrics vary significantly. We have found that there is no consensus on what trust means and that there is no study that has approached trust in utilizing a combination of discretionary and mandatory trust policies.

We explored the development of trust in the use of information technology, network information systems, and trust in relationships. Trust in relationships sounds simple but it is not. This is because trust in relationships is dynamic (trust changes as the relationship develops or changes). This is to also include trust from relationships from other countries, for example, between U.S. and Iran, or foreign governments or militaries. Trust is the “glue” of the relationship: It facilitates authenticity and respect. We also covered the process of protecting sensitive information and how a security clearance and a need-to-know policy can be applied to trust. We also explored the use of trust in digital signatures and how they are being applied to e-commerce to authenticate the identities of the parties’ transactions

We explored different authors' definitions of trust before we attempted to build a computational model of trust. It appears that in some situations organizations attempt to address the concept of trust, without ever defining trust. Without a formal and commonly accepted definition and identification of the components of trust, how can an organization effectively deal with the issues related to trust? We also discussed security issues, some simple applications to computer security, and we explored the notion of trust in commercial-off-the-shelf products.

We explored some of the basic concepts that appear in formal models. We showed that secure computing systems may be decomposed into data structures, processes, information about users, input-output (I/O) devices, and security attributes for controlled entities. The primary aspects of models include policy objective, locus of policy enforcement, strength of policy enforcement, granularity of user designations, and the locus of administrative authority. This was key in the development of the D-M model since it incorporates both discretionary and mandatory policies.

We introduced the D-M Trust Model, which provides for representing both discretionary and mandatory policy about trust; the D-M model of trust, addresses the need to model both of these types of policies explicitly and concurrently. The D-M model is built on top of, or is intended to be used in conjunction with, existing computational models of trust.

We do not claim the work on the D-M model is definitive or conclusive in its current state, but rather that it is a step forward in the understanding of the intricacies of discretionary and mandatory trust. Of course, more work is needed. In particular, we

believe that there are different levels of abstraction, which possibly implies the need for multiple levels of models. There is a time aspect to the dimension of trust in models one should consider. We categorize these as: strategic, operational, and tactical. The strategic aspect of the time line focuses on the long term. The focus in the operational dimensions is on the medium term, while the tactical dimension is focused on the short term.

We provided examples of how to apply the D-M model. In the first case study, we model how both discretionary and mandatory policies about trust affect the entire Aegis class ship, from the lowest actor, to the most senior actor. In our second case study of Battle Group Connectivity, we modeled trust along multiple dimensions, including top-down, bottom-up, and lateral flow of both information, multiple levels of trust-based decision making, and trust-based interactions among multiple actors who can form and dissolve actor-to-actor relationships on the fly.

B. FUTURE WORK

1. Public Key Infrastructure

One way to implement trust in B2B, B2G, or I2I, is by implementing a public key infrastructure (PKI). PKI enables users of a non-secure public network such as the Internet to securely and privately exchange data and money through the use of a public and a private cryptographic key pair that is obtained and shared through a trusted authority. The public key infrastructure provides for digital certificates that can identify individuals or organizations and directory services that can store and, when necessary,

revoke them. Although the components of a PKI are generally understood, a number of different vendor approaches and services are emerging.

The public key infrastructure assumes the use of public key cryptography, which is the most common method on the Internet for authentication a message sender or encryption of messages. Traditional cryptography usually involves the creation and sharing of a secret key for the encryption and decryption of messages. The secret or private key system's Achilles heel is the key; if it is discovered or intercepted by someone else, messages encrypted with that key can be decrypted. For this reason, the D-M model would be useful for identifying trust relationships in the development and maintenance of PKI. Applied to public key cryptography and the public key infrastructure, the D-M model would create a more trusted environment and reliable infrastructure to conduct data exchange on the Internet.

2. Influence Net Modeling

Situational Influence Assessment Modeling (SIAM) is a software application tool designed to allow the user to construct and analyze complex influence net models. SIAM claims it can help organize and evaluate large amounts of potentially conflicting information, facilitate richer analysis, and simplify complex decision making. Its key features include: 1. Automated assessment – allows users to evaluate complex paths of influence and adjust decisions to incorporate new observations, as well as run alternative projections of the events most likely to change the situation. 2. Documentation – enables users to include supporting sources, including rationales, citations to relevant information, interviews, and other resource materials. 3. Centralization – helps

coordinate functions typically performed by multiple experts, who may contribute to the decision making role. 4. Export Capabilities – allows users to export their model to a Bayesian Interchange Format file for import to any compatible software application.

(Influence Net Brochure)

The Naval Postgraduate School is fortunate to have Situational Influence Assessment Modeling (SIAM) software installed in a secure lab. The authors feel that one could easily take the data articulated in this thesis and introduce it to SIAM. As a result, one would be capable of generating measures of effectiveness for organizational trust through the use of SIAM. However, the SIAM model is only as good as the users assumptions.

3. Decision Support

Another recommendation of future work is to incorporate the D-M model into one of many decision support systems. The decision support methodologies represents a portion of the current state of management science's efforts to assist decision-makers in solving multi-criteria decision problems. Decision support methodologies are well founded, in theory and in practice, and have existed for over 25 years. These decision support methodologies and their software applications, like all other decision support tools, are neither designed to replace the decision-maker nor diminish the responsibility for the decision made. Both of them are capable of representing a non-trivial decision process in an effort to expand the rational boundaries of those involved in the decision-making process. These and other decision support tools only serve to prompt an analysis of as much of the available information as the decision-maker desires.

THIS PAGE INTENTIONALLY LEFT BLANK

APPENDIX. GLOSSARY

Address Spoofing: Altering the TCP/IP packets to make it appear that the message came from a source other than the originator.

Authentication: The process used to ascertain the identity of a subject.

Availability: Ensures that computer assets are fully operational when needed.

Back Door: An undocumented access code or procedure for accessing information.

Bridge CA: Instead of the CAs cross certifying with each other, they cross certify with a third party "bridge CA" that acts as an intermediary between CAs.

Certificate: A data structure that securely links an entity with its corresponding public key.

Certification Authority (CA): The component of the public key infrastructure that is responsible for issuing, revoking and certifying public keys.

Certificate Revocation List (CRL): A list of certificates that have been cancelled before their expiration date.

Ciphertext: The output of an encryption algorithm, or the encrypted form of a message.

Confidentiality: Ensures that information within a computer or transmitted can only be read by authorized personnel.

Cryptography: The branch of cryptology that deals with the algorithms that encrypt and decrypt messages or files to provide security and/or authenticity. (Stallings, W., 1999)

Digital Signature: An authentication mechanism that utilizes public key cryptography to guarantee the source and integrity of a message.

Domain: The logical realm over which a CA determines policy.

FORTEZZA: "FORTEZZA®" is a registered trademark held by the National Security Agency. It describes a family of security products. The FORTEZZA crypto card started as a low cost security device for the Defense Message System. However, the card was designed to be a general purpose cryptographic "co-processor" that can be used in numerous applications. The DoD class 4 PKI system uses FORTEZZA standards.

Hackers: People who abuse information systems or use them to commit criminal acts.

Hash Function: A function that combines a bit string with a secret key to generate a fingerprint of the message. The recipient of the message uses the same key to generate a hash value of the message and compares the two hash values. If they are the same, the message's integrity is valid.

Integrity: Only authorized personnel can modify computer assets or transmissions.

Key: A string of bits used in encryption algorithms to encrypt plaintext and decrypt ciphertext. The string's length depends upon the type of algorithm used.

Local Registration Authority (LRA): The person or organization that is responsible to be CA for properly identifying an entity seeking a certificate.

Lightweight Directory Access Protocol (LDAP): The defacto standard for accessing directory systems.

Nonce: An identifier or number that is used with authentication techniques to combat the man-in-the-middle attack.

Non-Repudiation: A message is sent such that the identity of the sender and the integrity of the message are strong enough to prevent that party from later denying that the transaction ever occurred.

Plaintext: The message that is to be encrypted, or the message that is recovered from decryption.

Pretty Good Privacy (PGP): A public-key cryptography program that was developed primarily by Phil Zimmerman in 1991.

Private Key: One of two keys used in public key cryptography. The private key is known only to the user and should be kept secret. Only the user should have the private key. The private key decrypts the corresponding public key.

Public Key: One of two keys used in public key cryptography. The public key is made available to everyone. The public key can decrypt its corresponding private key to verify authenticity (digital signature).

Public Key Cryptography: Cryptography that uses a pair of related keys to perform cryptography. When the keys are generated, one is designated the “private key”, which is kept secret and the other key is the “public key”, which is available to everyone. Public key cryptography is also called asymmetric cryptography.

Public Key Infrastructure (PKI): The key management system that ensures public keys are safely, efficiently, and conveniently delivered to the system that needs them.

Registration Authority (RA): In many cases the actual identity verification is delegated from the CA to another organization called registration authority (RA).

Root Certificate Authority: The most trusted entity in a hierarchical PKI domain. It is responsible for establishing and maintaining the PKI domain. It establishes the policy, issues the certificates and delegates responsibilities to lower level CAs or LRAs. It is the trust anchor.

Subjective: The evaluation of an object or occurrence is unique to each person.

Subjective Logic: It consists of a set of algebraic operators. It can be called a calculus for uncertain probabilities.

Symmetric Cryptography: The same key that is used to encrypt the message is used to decrypt the message.

Transitivity: In the context of trust, in order for trust to be transitive in a trust path, trust must be valid for each member in the path. For example, Bob trusts Sue, and Sue trusts Tom, transitivity assumes that Bob trust Tom.

Trojan Horse: An innocent looking program that has additional malicious functions.

Trust Anchor: The CA that is fully trusted by a user. This means that the user has complete trust in the CA's public key.

Trust Models: They attempt to automate the logic, variables, and thought processes that a human performs when making a trust decision.

Trusted Path: The verification path that a user must take to verify a certificate with a trusted CA.

Virus: A self-replicating computer program. A virus is often malicious code embedded in an executable program.

Worm: A self-replicating program, but unlike a virus it does not need a host to propagate, it is designed to spread on its own. It is malicious in that it performs a denial of service attack.

X.509 Standard: The standard that defines the structure and functionality for certificates and CRLs.

THIS PAGE INTENTIONALLY LEFT BLANK

LIST OF REFERENCES

- Abdul-Rahman, A., and Hailes, S., "A Distributed Trust Model," NSPW '97. Proceedings of the Workshop on New Security Paradigms Workshop, pp. 48-60, 1997.
- Adams, C. and Llody S., "Understanding the Public-Key Infrastructure" January 2000.
- Barton, D., Moran, A., and O'Connor, L., "Design Issues in a Public Key Infrastructure (PKI)," [<http://www.csu.edu.au/special/auugwww96/proceedings/barmoroco/Barmoroco.html>], August, 1996.
- Bhimani, A., "PKI, Be careful what you wish for...", Information Security, May 2000, pp. 38-50.
- Booker, R., "Practical PKI," Messaging Magazine, [http://www.ema.org/html/pubs/mmv5n5/Practical_PKI.htm], September/October 1999.
- Briney, A., "PKIs: From Pilot to Production," Information Security, May 2000, pp. 54-60.
- Cabletron Systems., "Public Key Infrastructure (PKI)," [<http://www.cabletron.com/vpn/VPNpki.htm>], June, 1999.
- Cabletron Systems., "Public Key Infrastructure (PKI)," [<http://www.cabletron.com/vpn/VPNpki.htm>], June, 1999.
- Chu, Y., "Trust Management for the World Wide Web," Master's Thesis, Massachusetts Institute of Technology, Boston, Massachusetts, 13 June, 1997.
- Ford, W., "Public-Key Infrastructure Interoperation: Some Pragmatics," Messaging Magazine, [<http://www.ema.org/html/pubs/mmv3n5/pubkey.htm>], September/October 1997.
- Gaines, L., "Trust and its ramifications on the DoD Public Key Infrastructure," NPS Thesis, September 2000.
- Gerck, E., "Modelling Trust," [<http://www.sandelman.Ottawa.on.ca/spki/html/1998/winter/msg00077.html>], January, 1998.
- Gerck, E., "Towards Real-World Models of Trust: Reliance on Received Information," [<http://www.mcg.org.br/trustdef.htm>], January 1998.

Grant, Gail, "Understanding Digital Signatures: Establishing trust over the Internet and other Networks" April 2000.

Grogan, Janis, "Digital Identity Crisis," Information Week, 2 November 1999, pg 154.

Hansen, A., "Public Key Infrastructure (PKI) Interoperability: A Security Services Approach to Support Transfer of Trust," NPS Thesis, September 1999.

Hardin R., "Trusting persons, trusting institutions. In Strategy and Choice," Cambridge, MA: MIT, 1991 pp487.

Jøsang, A., "A Metric for Trusted Systems," Proceedings of the 21st National Security Conference, NSA, 1998.

Jøsang, A., "Modelling Trust in Informatio Security, PhD Research Project,1998," Abstract, [<http://www.item.ntnu.no/~ajos/PhD.html>], 1998.

Jøsang, A., "A Subjective Metric of Authentication," Proceedings of the 5th European Symposium on Research in Computer Security (ESORICS'98), Springer-Verlag, 1998.

Jøsang, A., "An Algebra for Assessing Trust in Certification Chains," Proceedings of the Network and Distributed Systems Security (NDSS'99) Symposium, The Internet Society, 1999.

Jøsang, A., "Artificial Reasoning with Subjective Logic," Proceedings of the Second Australian Workshop on Commonsense Reasoning, 1997.

Jøsang, A., "Trust-Based Decision Making for Electronic Transactions," Proceedings of the Fourth Nordic Workshop on Secure Computer Systems (NORDSEC'99), Stockholm, 1999.

Jøsang, A., "Prospectives for Modeling Trust in Information Security," Proceedings of the 1997 Australasian Conference on Information Security and Privacy, Springer, 1997.

Joint Vision 2020, Coordination Draft, 18 Feb 20000 Version.

Khare, R., and Rifkin, A., "Trust Management on the World Wide Web," [http://www.firstmonday.dk/issue3_6/khare/], June 1998.

Khare, R., and Rifkin, A., "Weaving a Web of Trust," [<http://www.cs.caltech.edu/~adam/local/trust.html>], 30 November 1997.

Kramer, R., "TRUST AND DISTRUST IN ORGANIZATIONS: Emerging Perspectives, Enduring Questions," January 1, 1999, Annual Review of Psychology, Pg. 569.

Lubbe, J., "Basic Methods of Cryptography," Cambridge University Press, 1998, Pg. 204.

Menezes, A., Oorschot, P., and Vanstone, S., "Handbook of Applied Cryptography," CRC Press, 1997.

Myers, A., and Liskov, B., "A Decentralized Model for Information Flow" ACM SIGOPS Operating System Review, Vol. 3, No. 5, December 1997, pp. 129-142.

Maccoby, Michael, "Building Trust Is An Art," Research-Technology Management, October 1997, Vol. 40, No. 5, Pg. 56-57.

McKnight, D. Harrison, and Chervany, Norman, "The Meaning of Trust," MISRC Working Paper Series, [<http://www.misrc.umn.edu/wpaper/wp96-04.htm>], April 1994.

Rapid Anti-Ship Cruise Missile Integrated Defense System (RAIDS), [<http://www.fas.org/man/dod-101/sys/ship/weaps/an-syq-17.html>], December 1998.

Ship Self-Defense System, SSDS, [<http://www.ratheon.com/products.html>], February 2000.

Sibley, E., Michael, J., and Wexelblat, R., "Use of an Experimental Policy Workbench: Description and Preliminary Results." C.E. Landwehr and S. Jajodia, eds., Database Security, V: Status and Prospects. Elsevier Science Publishers, Amsterdam, Neth., March 1992, pp. 47-76.

Toffler, Alvin, "Third Wave," Reissue edition, December 1991.

Reiter, M., and Stubblebine, S., "Authentication Metric Analysis and Design," in ACM Transactions on Information and System Security, Vol. 2, No. 2., May 1999, pp. 138-158.

Wyatt, Donna, "Trust is power," Executive Excellence, Shelton Marketing Communications 1996, Vol. 13, No. 12 Pg. 12-13.

THIS PAGE INTENTIONALLY LEFT BLANK

INITIAL DISTRIBUTION LIST

1. Defense Technical Information Center 2
8725 John J. Kingman Road, Ste 0944
Fort Belvoir, VA 22060-6218
2. Dudley Knox Library 2
Naval Postgraduate School
411 Dyer Road
Monterey, CA 93943-5101
3. Professor Carl R. Jones 1
Code IS/JS
Naval Postgraduate School
Monterey, CA 93943-5118
4. Professor James Bret Michael 1
Code CS/Mj
Naval Postgraduate School
Monterey, CA 93943-5118
5. Professor Audun Jøsang 1
DSTC Pty Ltd
Level 7, GP South (Bldg 78)
The University of Queensland
Brisbane, QLD 4072
Australia
6. LCDR Leonard T. Gaines 1
Naval Supply Systems Command (Code 63G)
Bldg 309, Room 305
5450 Carlisle Pike
Mechanicsburg PA, 17055-0791
7. LCDR Anthony Hansen 1
Naval Research Laboratory
Bldg. 259/Code 9110
4555 Overlook Ave.
Washington, D.C. 20375

8.	Mr. Terry Mayfield	1
	Computer and Software Engineering Division	
	Institute for Defense Analysis	
	1801 North Beauregard Street	
	Alexandria, VA 22311-1772	
9.	Professor John McHugh	1
	SEI/CERT	
	4500 5 th Avenue	
	Room 4420	
	Pittsburgh, PA 15213-3890	
10.	LT Carl M. Pedersen	2
	26133 Legends Court	
	Salinas, CA 93908	